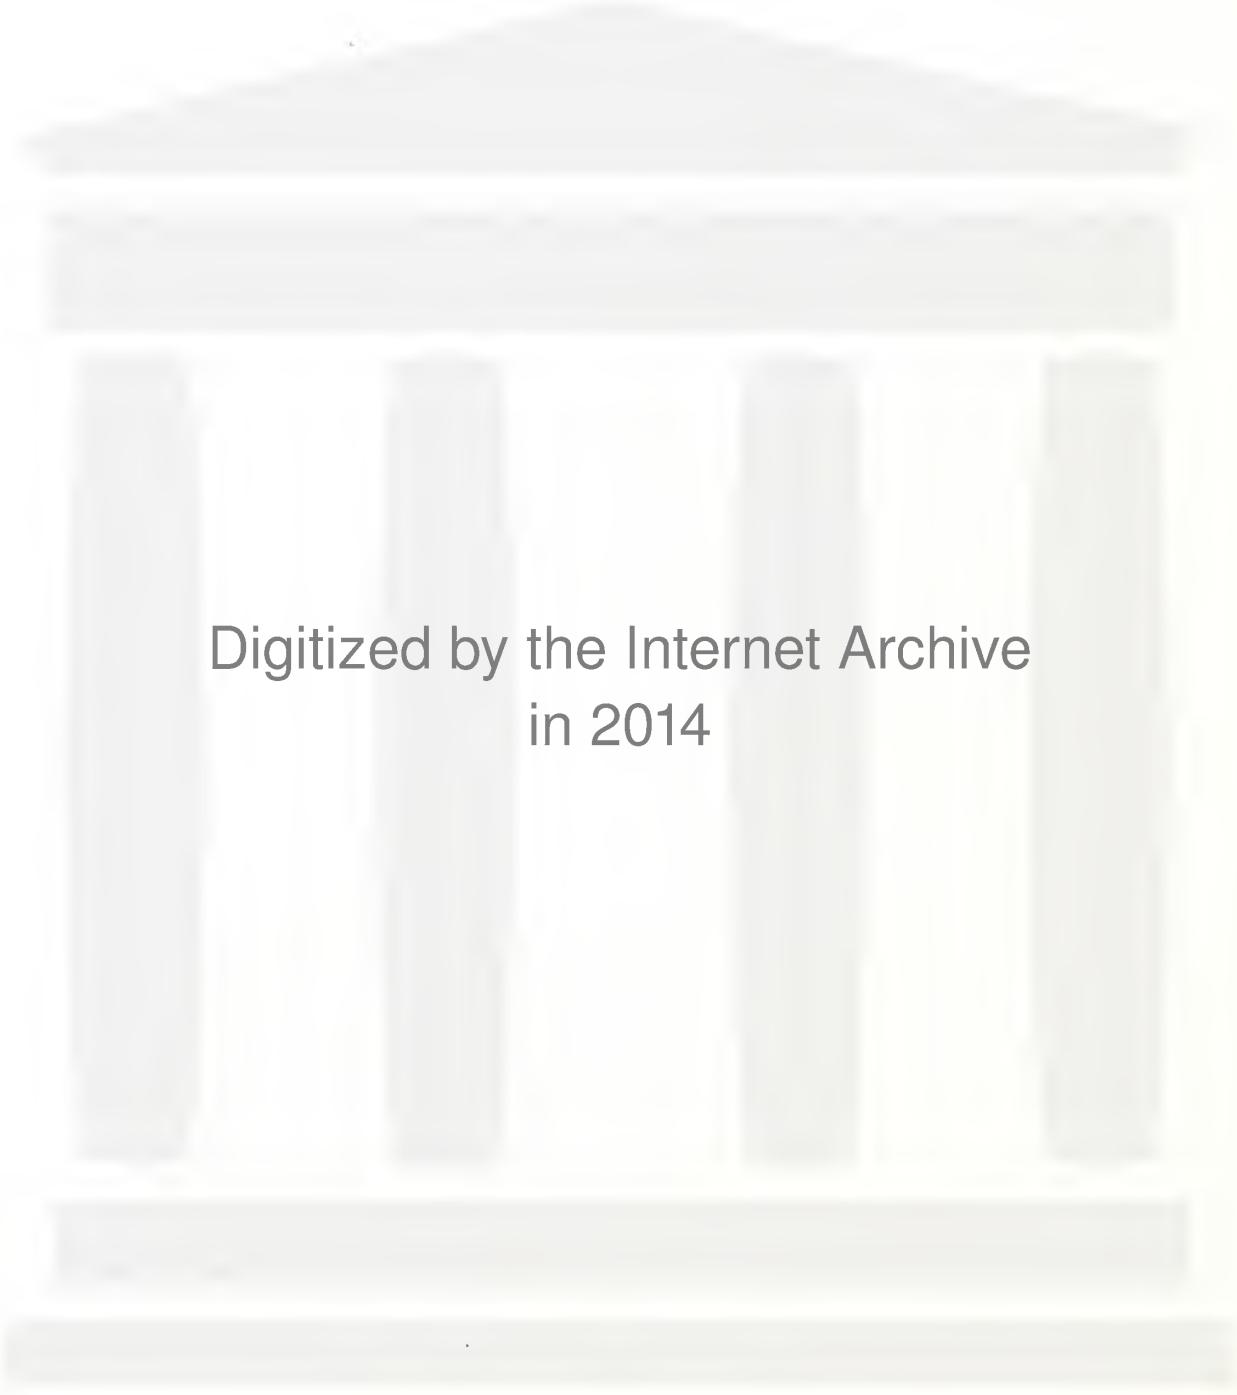


ISP

**Large-Scale Systems
Directions - Large IBM
and Software-
Compatible Mainframes**



Digitized by the Internet Archive
in 2014

<https://archive.org/details/20220MSDBxx84LargeScaleSy>

LARGE-SCALE SYSTEMS DIRECTIONS LARGE IBM AND SOFTWARE COMPATIBLE MAINFRAMES

U-CLS3
1984 c.1

AUTHOR Large-Scale Systems Directions--

TITLE

IBM & S/W Compatible Mainframes

第八章

IBM & S/W Compatible Mainframes

DATE
LOANED

DATE LOANED	BORROWER'S NAME
----------------	-----------------

~~5 | 3 Brussels
2 | 7 Tech~~





**LARGE-SCALE SYSTEMS DIRECTIONS
LARGE IBM AND SOFTWARE-COMPATIBLE MAINFRAMES**

CONTENTS

	<u>Page</u>
I INTRODUCTION	I
II GENERAL SYSTEMS TRENDS AND MAINFRAME ARCHITECTURES	3
A. General Systems Concepts	3
B. Mainframe Centralization of Function	4
1. The von Neumann Architecture	4
2. IBM Operating Systems	7
C. Integration to Infinity?	8
1. SNA--The Host with the Most	8
2. MVS/XA and Multiprocessing	9
3. The Problems of Office Automation	10
4. The "Virtual Bottleneck"	12
D. Differentiation of Function	13
1. The Current Architectural Environment	13
2. Vector, Array, and Matrix Processors	15
3. Data Base Computers	20
4. Network Management	23
E. Mechanization and Micros	24
III RESIDUAL VALUE FORECASTS	27
A. Factors Affecting Residual Value Forecasts	27
B. Announcements	30
C. Projected Used Market Prices and Residual Values	34

**LARGE-SCALE SYSTEMS DIRECTIONS
LARGE IBM AND SOFTWARE-COMPATIBLE MAINFRAMES**

EXHIBITS

		<u>Page</u>
II	-1 IBM Centralization and Integration	14
	-2 An Example of Architectural Differentiation	19
	-3 Sample Processing Configurations	21
III	-1 Factors Affecting Computer Equipment Residual Values	28
	-2 Projected Used Market Retail Values	35
	-3 Projected Residual Values	37
	-4 Residual Value Forecast for IBM 4331-K02 Processor	39
	-5 Residual Value Forecast for IBM 4341-L02 Processor	40
	-6 Residual Value Forecast for IBM 4361-K0 Processor	41
	-7 Residual Value Forecast for IBM 4381-K05 Processor	42
	-8 Residual Value Forecast for IBM 3083-EX8 Processor	43
	-9 Residual Value Forecast for IBM 3083-BX16 Processor	44
	-10 Residual Value Forecast for IBM 3083-JX32 Processor	45
	-11 Residual Value Forecast for IBM 3081-GX16 Processor	46
	-12 Residual Value Forecast for IBM 3081-KX24 Processor	47
	-13 Residual Value Forecast for IBM 3084-QX64 Processor	48
	-14 Residual Value Forecast for Amdahl 5850-24 Processor	49
	-15 Residual Value Forecast for Amdahl 5860-24 Processor	50
	-16 Residual Value Forecast for Amdahl 5868-32 Processor	51
	-17 Residual Value Forecast for Amdahl 5870-32 Processor	52
	-18 Residual Value Forecast for Amdahl 5880-64 Processor	53
	-19 Residual Value Forecast for NAS AS/6630	54
	-20 Residual Value Forecast for NAS AS/6660	55
	-21 Residual Value Forecast for NAS AS/8023	56
	-22 Residual Value Forecast for NAS AS/8083	57
	-23 Residual Value Forecast for NAS AS/9050	58
	-24 Residual Value Forecast for NAS AS/9070	59

I INTRODUCTION

- In 1984, INPUT changed the title of its residual value series to Large-Scale Systems Directions, and this report completes the first set of the new series. The selection of the title was fortuitous, but it has also proved to be fortunate. In attempting to determine IBM's large-scale systems directions, one is filled with the same feeling of awe that is inspired when viewing the expanding universe from this minor planet. During the past year, IBM literally seems to have been expanding in all directions at the same time. This presented INPUT with the problem of establishing direction (or at least some order) out of chaos.
- The direction problem became especially difficult when attempting to determine IBM software strategies, and it was decided that the concepts of general systems theory (GST) as defined by Ludwig von Bertalanffy would serve as an appropriate means of relating IBM's directions to the inevitable trends of complex computer/communications networks. This methodology was defined and used in INPUT's reports: Market Impacts of IBM Software Strategies and Information Systems Implications of IBM Software Strategies.
- This report uses the terminology of general systems theory to describe trends in mainframe architecture, and Chapter II gives a general explanation of GST concepts. Essentially, the prevailing von Neumann architecture is analyzed, and the emergence of alternatives is discussed. These alternatives are currently manifesting themselves in the form of supercomputers and data base machines, but they are also immediate threats to the role of the large-scale, general purpose mainframe and its role as defined by IBM systems software.

- Chapter III presents updates of IBM and software-compatible mainframe residual value forecasts, which were contained in:
 - Residual Value Forecasts for Large-Scale Systems, December 1983.
 - Large-Scale Systems Directions: Midyear Update, August 1984.
- In addition, the general methodology used in forecasting residual values is reviewed in Chapter III. It is important that the methodology be understood when employing the forecasts.

II GENERAL SYSTEMS TRENDS AND MAINFRAME ARCHITECTURES

A. GENERAL SYSTEMS CONCEPTS

- Earlier this year, confronted with the problem of analyzing IBM software strategies, INPUT found it convenient to apply the concepts of general systems theory. (Market Impacts of IBM Software Strategies and Information Systems Implications of IBM Software Strategies, INPUT, 1984.) Essentially, the theory states that in any complex system progressive centralization, integration, differentiation, and mechanization occur in parallel. These relatively simple concepts are defined as follows:
 - Progressive centralization: "Leading parts" develop and dominate the behavior of the system.
 - Progressive integration: The parts become more dependent upon the whole.
 - Progressive differentiation: The parts become more specialized.
 - Progressive mechanization: Parts become limited to a single function.
- Although the concepts are relatively simple, they provide a convenient framework to view and analyze hierarchical systems of extreme complexity. More importantly, it is INPUT's opinion that an understanding of these funda-

mental concepts facilitates the prediction of technological directions in computer systems and the identification of "abnormalities" that may arise from undue emphasis upon the GST trend at the expense of others. For example, early attempts at standardization of programming languages (a form of mechanization) were doomed to failure because they ignored the nature (and inevitable) trend toward differentiation as a broader user and applications base developed.

- PL/I represented an attempt by a vendor (IBM) to establish a single programming language (function) as a standard. Although this attempt failed, the parallel effort to establish a mainframe hardware standard has been eminently successful. As we all await the next large-scale mainframe offering (Sierra) from IBM, it is becoming increasingly apparent that the IBM mainframe architecture is less than divinely inspired and may actually be an "abnormality" with which we have all learned to live.

B. MAINFRAME CENTRALIZATION OF FUNCTION

I. THE VON NEUMANN ARCHITECTURE

- 1984 is the fortieth anniversary of the "von Neumann machine," which fundamentally treats its programs as if they are data in storage. This magnificent conceptual breakthrough permitted the development of general purpose computers as we know them today. Prior to that, instructions were hand-wired into the processor.
- The general purpose computer provides the "leading part" that can control both the flow of processing and the flow of data. Conceptually, it facilitates centralization.

- Unfortunately, the general purpose computer also creates what John Backus (the father of FORTRAN) has described as the "von Neumann bottleneck." Treating instructions and data in a similar fashion means that data are transferred between storage and the processor one word at a time. Despite the tremendous advances in processor speeds and memory sizes on large general purpose mainframes, these machines still funnel data (including instructions) back and forth through this narrow bottleneck.
- During the past year, INPUT has attempted to put large-scale mainframes into perspective.
 - Their primary functions were listed in the Residual Value Forecasts for Large-Scale Systems, December 1983. These functions were defined as:
 - Heavy computation.
 - Transaction processing against large data bases.
 - Data base management for a computer/communications network.
 - In Large-Scale Systems Directions: Disk, Tape, and Printer Systems, published in March 1984, INPUT projected IBM's probable structuring of distributed data bases. On large-scale mainframes these data structures were:
 - Physical sequential files.
 - VSAM files.
 - IMS/VS and DL/I data bases.
 - DB2 relational tables.

- In Large-Scale Systems Directions: Midyear Update, INPUT warned of potential large mainframe performance problems based upon the use of large general purpose mainframes as "data base machines." Some of these performance problems can be directly related to IBM's highly centralized hardware/software implementation of large general purpose mainframes of von Neumann architecture.
- Specifically, the von Neumann architecture is designed for scalar operations; and both scientific and commercial computer operations are dealing with vectors, lists, tables, and strings where one word at a time does indeed represent a severe bottleneck. In addition, the IBM System/360 and 370 implementation of von Neumann machine leaves a lot to be desired in the handling of such structures. For example, even before System/360 was announced in 1964, computers with the following features were available:
 - Multiple levels of automatic indirect addressing and address modification, which minimized processor access to storage during instruction interpretation and eliminated unnecessary movement of data between the processor and storage (for example, during sorting).
 - Automatic memory protection through the setting of limit registers and interrupts on tagged data and instruction words (eliminating the looping required to check table entries, record blocking, and interference between programs and data).
 - Word sizes, which permitted more comprehensive addressing schemes (and precision of computation) without resorting to architectural changes such as XA (which has been referred to as "extended accommodation" within IBM). In other words, the 32-bit word size of the IBM 360/370 architecture tended to make the bottleneck more confining.

2. IBM OPERATING SYSTEMS

- SNA is 10 years old, but the large IBM host under MVS/XA has continued to epitomize progressive centralization as it strives to maintain its dominance of the network. Precious little function has been distributed from large mainframes over the past 10 years, and IBM has been reported as saying that MVS/XA will "run out of gas" by the late 1980s because it will not be able to handle all of the intelligent workstations that will be linked up by that time.
- Think about the large central host in an SNA network controlling hundreds of intelligent workstations (soon to be thousands), tens of gigabytes of direct access storage (soon to be hundreds of gigabytes or terabytes), tens of megabytes of RAM (soon to be hundreds), and the variety of tasks being performed under MVS/XA, TSO, IMS, DB2, etc. Then think about programs (systems and applications) designed for the von Neumann architecture being passed (along with their data) through the bottleneck between processor and storage. It is a miracle that these centralized systems have progressed to their current size and work at all!
- Yet they remain at the heart of IBM's hardware/software strategy through the 1980s; and despite reports to the contrary, the large IBM mainframes (and associated operating systems) are not going to become extinct because of either new hardware (microprocessor) or systems software (UNIX) advances. However, there is no question that IBM has emphasized progressive centralization at the expense of other GST trends, and it is primarily a question of how much control IBM can exercise over the inevitable trends--toward integration, differentiation, and mechanization--inherent in computer/communication networks.

C. INTEGRATION TO INFINITY?

I. SNA--THE HOST WITH THE MOST

- SNA originally was designed to exercise control over distributed processing by employing relatively expensive, limited function controllers (3705, 3790) in lieu of minicomputers, which were more cost-effective for distributed processing. Any processing at remote locations would be integrated by becoming host dependent. For example:
 - Timesharing from outside service bureaus and departmental minicomputers was countered by providing a timesharing option (TSO) under what was essentially a batch system (OS/360, evolving over time to MVS/XA). However, individual terminals serviced directly from a central host were not cost-effective compared to interactive terminals serviced from either local or remote minicomputers.
 - When specific interactive applications started appearing in branches of banks, retail outlets, insurance companies, etc., it became apparent that the large number of terminals being serviced could not be supported from remote large-scale mainframes. IBM's initial response was to distribute a few communication functions to the 3705 and minimal processing to the 3790 cluster controller and keep all data bases (and software) centrally located.
 - Eventually some software was developed (DMS was a joint IBM-customer software project) and it became apparent that even at best the 3790 would not support any significant amount of processing. The economics of minicomputers (including IBM's own Series/1) became increasingly apparent, and IBM's response was (and continues to be) the excruciatingly slow upgrading of the 3705 to the 372X, and the 3790 to the 8100 series.

- The demand for large mainframe MIPS was fueled by providing terminal access to centralized data bases--especially when the DBMS was IMS.
- The unexpected, rapid development of personal computers (microprocessors) is causing forced integration under the great SNA umbrella. Micro-mainframe links are designed to make the standalone PC (which has already been semi-officially declared dead by IBM) become dependent upon the host for data.
- There are also a multitude of standalone processors (System 36s, 38s, 43XXs, Series/1s, and foreign processors) that must eventually be incorporated in the network at least at the level of document interchange and control, and for data base transfer and backup. Integration under SNA will result in additional host burden.

2. MVS/XA AND MULTIPROCESSING

- MVS/XA is the latest extension of IBM's mainstream operating systems thrust, which started with OS/MFT and has gone through numerous iterations. However, it is not an appropriate system to integrate all of the various systems mentioned above. Therefore, it will probably find itself integrated under VM along with a variety of operating systems (including UNIX) at various levels in the processing hierarchy. Piling another level on the software hierarchy only places additional burden on the mainframe.
- As an architectural answer to the demand for more MIPS, dyadic processors were employed in the 3081, and the 3084 is an MP version of the 3081. There is overhead associated with multiprocessing systems (see Exhibit II-1 in Large-Scale Systems Directions: Midyear Update), and integrating a variety of networked systems through a centralized mainframe with quadruple processors running multiple operating systems under VM is going to result in one big bottleneck in the SNA network.

- However, this massive centralization seems to be the way IBM is going, because Sierra is being promised with a "single systems view." Indeed, it is INPUT's opinion that integration of distributed data bases requires the emergence of a "leading part" to dominate the behavior of the network; otherwise, chaos is inevitable. Unfortunately, IBM's traditional emphasis upon progressive centralization (without orderly distribution of processing to its proper place in the network) has already resulted in an abnormal burden on host mainframes (many large installations are growing at the rate of 50% per year in required processing power). Therefore, data base integration through the SNA hosts may result in the necessity to distribute processing power in a less than orderly fashion with resulting negative impacts on information quality.

3. THE PROBLEMS OF OFFICE AUTOMATION

- Emerging office systems that integrate voice, data, text, and images present especially thorny problems for IBM under their heavily centralized, host-oriented strategy. Once paper starts to disappear (or at least be contained) and electronic mail and messages are substituted for some voice communications, a new level of reliability/availability/serviceability is required. And, regardless of improvements that have been made, IBM mainframe hardware/software RAS does not compare favorably with, let's say, an electronic switching system for AT&T communications.
- In addition, performance of general purpose hardware/software systems, in terms of the number and cost per terminal serviced, has never compared favorably with systems designed specifically for the communications environment. Theoretically, the operator at the intelligent workstation should not have to be concerned about where data or information is located, but there is going to be a big difference in responsiveness when it is located on the LAN and when it is on the mainframe (even if the mainframe is in the same building).

- The fact that a request forwarded to the mainframe may require searching an enormous data base will be immaterial to the new class of users. Consistency of response is extremely important in an interactive environment, and anything going through the mainframe bottleneck is going to be subject to delay.
 - There is also a question of cost and billing. Data and documents obtained from the central storage facility (mainframe) are going to be substantially more expensive than those stored, produced, and retrieved on the LAN.
 - It is little wonder that IBM is approaching LANs cautiously. The mainframe integration under SNA is going to expose price performance imbalance of substantial proportions.
- Office systems will also expose the technological weaknesses of the mainframes that support them. For example, integrated electronic filing systems incorporating optical disks are beginning to appear. Meanwhile, IBM is supporting Scanmaster I off of mainframes.
 - The cost of storing images on optical disk versus magnetic disk is going to be lower by at least an order of magnitude. (See Exhibit II-4 in Large-Scale Systems Directions: Disk, Tape, and Printer Systems, INPUT 1984.)
 - In addition, the cost of image processing on large mainframes is going to expose not only the "von Neumann bottleneck," but also the "virtual bottleneck" of page size. An MVS page won't hold a picture (image): it will require at least 10 of them.

4. THE "VIRTUAL BOTTLENECK"

- The overhead of virtual storage systems has been largely ignored in recent years, but there is little question that it has contributed to the insatiable demand for processing power. It is not our intention to rehash the arguments concerning the benefits and costs associated with VS except to state that both technology (in terms of large, cheap main memory) and the role of large central mainframes have changed considerably since VS was determined to be the preferred scheme for memory management.
 - Because many programs requiring heavy computation demonstrate "locality of reference" within their inner loops and within their arrays of data, the cost of paging is not nearly so important as the "von Neumann bottleneck" in such scientific computation.
 - However, commercial application requiring sorting, table searching, and random access to large data bases do not exhibit locality of reference in either their program structure or data accessing. The use of host mainframes as data base machines requires precisely this type of processing, and page size (and paging management) becomes a very real "virtual bottleneck."
- It does not appear that IBM intends to change its current 308X, VM/MVS/XA-oriented architecture in either Sierra or Summit. However, it is INPUT's opinion that rapid off-loading of the mainframe through differentiation and mechanization of function may prove necessary if the host is to fulfill its proper role in the emerging distributed data base environment.

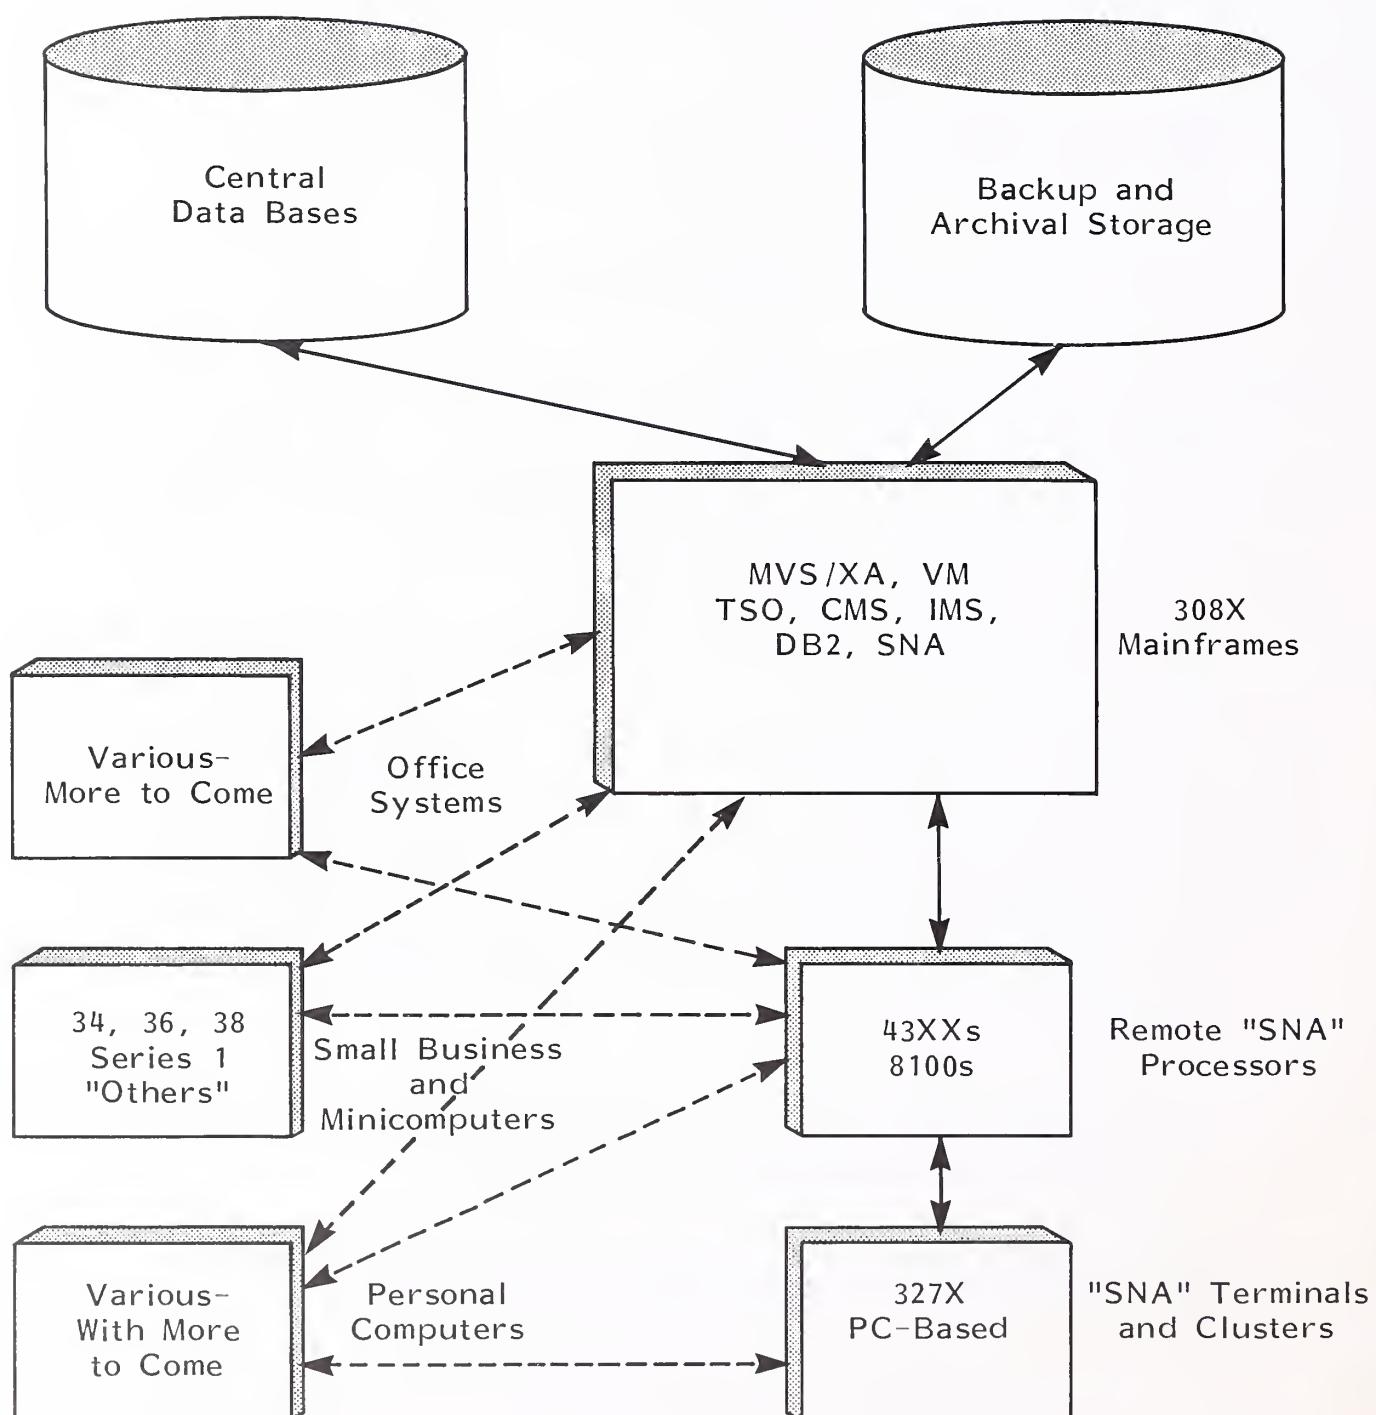
D. DIFFERENTIATION OF FUNCTION

I. THE CURRENT ARCHITECTURAL ENVIRONMENT

- Before proceeding to describe the inevitable GST trends toward differentiation and mechanization of function, it is important to recognize "abnormality" that has been created by IBM's past (and continued) reluctance to distribute processing on either an architectural or geographic basis. Exhibit II-1 depicts IBM's emphasis upon centralization and the pressures created by the current necessity to integrate personal computers, decentralized small business systems and minicomputers, and disparate office automation products.
- The IBM strategy is to get all of the sources of data and information integrated at some level under the highly centralized, host-oriented, SNA umbrella. The large mainframe becomes the "leading part" upon which all of the subsidiary parts become dependent for data and information interchange. (See Exhibit II-2 in Large-Scale Systems Directions: Midyear Update for the services provided by large host, data base machines.)
- The problem is that the host hardware/software architecture will create an expensive, central bottleneck that may not be capable of providing adequate service to the dependent parts of the system. When performance problems have arisen, the solution has been to dedicate systems based on applications: production versus development, batch versus interactive, batch/production versus data base, and now data processing center versus information center. From the vendors' point of view, this is an ideal "solution" for the following reasons:
 - Multiple systems are not as effectively used in terms of both processors and storage (more hardware sales).
 - Systems software must be duplicated (more software sales).

EXHIBIT II-1

IBM CENTRALIZATION AND INTEGRATION



- Performance problems (including reliability, availability, and serviceability) are obscured.
- Unfortunately, from the user's point of view the additional costs go above and beyond the direct hardware/software expense.
 - Each additional console (system) results in added costs in terms of overhead (flow space, lighting, air conditioning, plumbing, etc.), systems software personnel, operations personnel, and general expenses associated with coordination and management of multiple systems (even if they are on the same site).
 - Differentiating the users' applications systems by assigning them to hardware/software clones of the bottleneck is expensive, and it doesn't solve the problem. It merely creates the requirement for a new level of centralization and integration. The solution is to differentiate and mechanize the functions of the bottleneck.

2. VECTOR, ARRAY, AND MATRIX PROCESSORS

- The need to break the von Neumann bottleneck has long been recognized by those concerned with the development of supercomputers for heavy scientific computation and for many specialized applications in the federal government. In addition, the need to give a boost to the general purpose mainframe has long been recognized by the availability of attached scientific processors, which have been referred to as the "poor man's supercomputer." Even IBM has recognized and addressed this need with the IBM 3838.
- This year both Amdahl and NAS announced supercomputers that can be attached to, or be frontended by, their IBM software-compatible mainframes. Both companies anticipate "commercial" applications that will require heavy computation, but the systems differ substantially in their capabilities.

- The Amdahl 1100 and 1200 vector processors have rated speeds of up to 267 megaflops (millions of floating point instructions per second) and 533 megaflops respectively, and with purchase prices of \$9.2 million and \$13.7 million.
 - . These processors truly qualify as supercomputers and are obviously designed to compete in that marketplace.
 - . The ability to translate and subsequently optimize existing FORTRAN programs written for Amdahl (and IBM) mainframes implies that current applications of multiple Amdahl 58XX or IBM 308X systems might grow sufficiently to warrant such processing power. This appears doubtful, though, since the power differential is so great. (Megaflops on the vector processor cannot be compared with MIPS ratings on an IBM-compatible mainframe running under MVS: the applications throughput on the vector processor would be substantially greater than the megaflops/MIPS ratio would indicate.)
 - . The increased cost of the Amdahl 1100 and 1200 (maintenance alone runs \$35,900 and \$51,000 per month, respectively) precludes incremental growth from existing applications. The justification must be in new applications (or against existing supercomputers).
- NAS, on the other hand, announced the AS/9100 rated at 28 megaflops and selling for \$300,000-600,000 when added to NAS AS/9000 mainframes (the high end being for multiprocessor systems). The AS/9100 has been billed as an "entry-level supercomputer" that outperforms IBM 308X series uniprocessors 14-fold for vector processing.

- . At that price-performance level it is targeted more at current attached scientific processors such as the IBM 3838, which it is reported to outperform by approximately three to one for vector processing.
 - . Therefore, current FORTRAN workload could be adapted to run on the AS/9100 and could result in substantially better performance than on a machine of von Neumann architecture.
 - . There is a FORTRAN preprocessor that allows conventional scalar code to be translated to vector code "without programmer intervention," so it would theoretically be possible to differentiate scientific applications for running on the specialized processor.
 - . Depending upon the mix of workload, the AS/9100 provides a means of relieving the current and anticipated burden on large general purpose mainframes.
- Despite having translators for scalar FORTRAN, the general issue of languages and efficient compilers for vector processors (or array processors) has not really been addressed by either Amdahl or NAS. In addition, using an IBM-compatible, MVS-oriented mainframe to handle the scheduling of the vector processor makes it subject to a certain portion of the overhead and cost associated with the large, general purpose mainframe.
- In addition to the vector processors for backending IBM-compatible mainframes, it is worth mentioning that an additional level of differentiation is possible by isolating the basic "matrix-matrix multiplication" function on a matrix computer. During 1984, GuilTech Research Company (GRC) demonstrated the SC-533, a matrix computer capable of giga computation rates on fixed point operations (GOPs: 1,000,000,000 multiplies or adds per second)

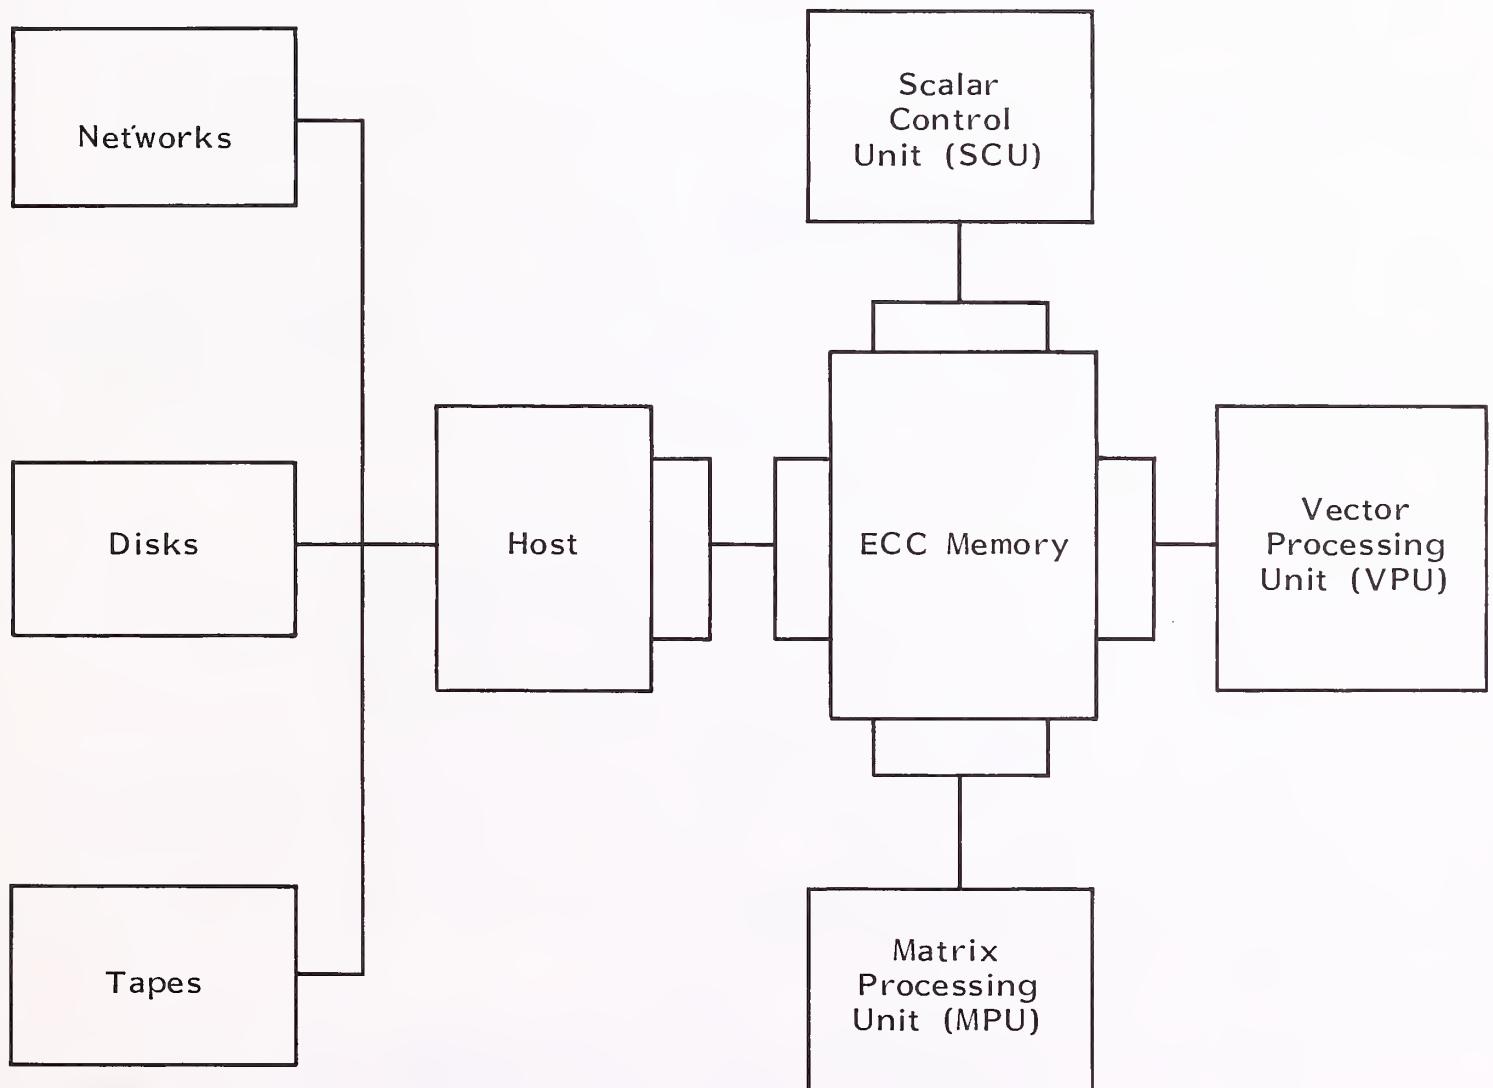
and 533 megaflops on floating point. The system is designed to interface to VAX, Cray, and Perkin-Elmer hosts; and to "networks" and "industry standard mass storage systems." Since matrix processing reduces the required data rate for a given computation, balance can be achieved between processing and I/O even at the data rates on minicomputer channels. The SC-533 is scheduled for delivery to beta-test sites in 1985 with production units available in 1986. The general architecture of the system is depicted in Exhibit II-2.

- The host serves primarily as a "data funnel" and not as a computational processor.
- The SCU (Scalar Control Unit) executes the user's applications program and, under control of the system executive, sequences and synchronizes the operation of the MPU (Matrix Processing Unit), VPU (Vector Processing Unit), and host I/O.
- The VPU is a standard commercial array processor (30 megaflops) that complements the matrix multiplication capabilities of the system with vector arithmetic necessary for many algorithms.
- The MPU performs the kernel operation of the matrix multiplication.
- All data and control communications (host to SCU, SCU to MPU, and SCU to VPU) occur via the memory, which is a minimum of 16 M-bytes, and all of the processing units operate in parallel.
- GRC intends to provide various software development features to permit user programs to be "matrixized" including an optional, off-line, software development tool for program development, debug, and simulation.
- The SC-533 represents the functional differentiation of an applications set (matrix-matrix multiplication), the architectural differentiation of

EXHIBIT II-2

AN EXAMPLE OF ARCHITECTURAL DIFFERENTIATION

GRC's SC-533



the hardware system (host, SCU, VPU, MPU), and the necessary software differentiation (both operational and development) to make effective use of the system.

- Therefore, a wide variety of "solutions" for the von Neumann computational bottleneck is becoming available. From the perspective of the large IBM host-oriented network, a variety of configurations is possible. A sample of these configurations is depicted in Exhibit II-3. But they have one thing in common: the general purpose mainframe serves as a data manager and processing scheduler for the specialized processors.

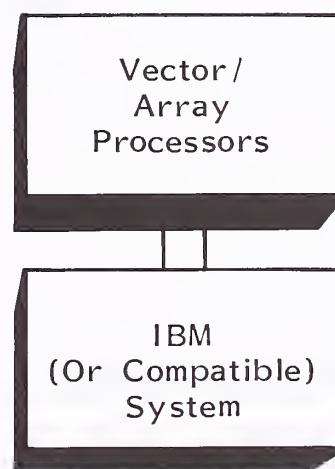
3. DATA BASE COMPUTERS

- Serving as a data manager for heavy computation is not the same as serving as a data base manager for large data bases. The volumes of data for scientific computation may be quite large, but normally they are reduced prior to heavy computation. Also, these data are normally represented by one- or two-word values (32 or 64 bits). The management of large data bases, on the other hand, results in complex structures of variable length data. The manipulation and control of these data structures also represent a substantial burden for the von Neumann architecture.
- Fortunately, many of the data management problems reduce to the processing of one- and two-word values in terms of indices, pointers, and sort keys. In addition, the structures of these data lend themselves to vector processing (regardless of whether they are called tables, lists, or directories) and eventually even to multidimensional matrix processing (for more complex data structures, including images). The primary difference is that data base management requires comparisons (logical operations) rather than conventional arithmetic.
- Unfortunately, while the need for fast arithmetic has been readily apparent in many scientific problems that simply can't be solved without ever-larger

EXHIBIT II-3

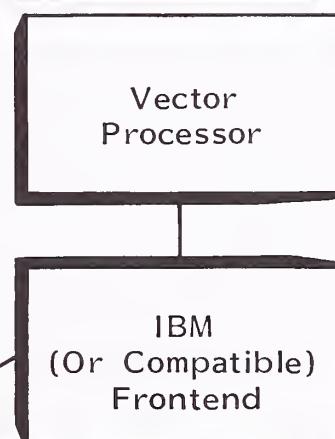
SAMPLE PROCESSING CONFIGURATIONS (Heavy Computation)

TIGHTLY COUPLED



IBM 3838, NAS AS/9100

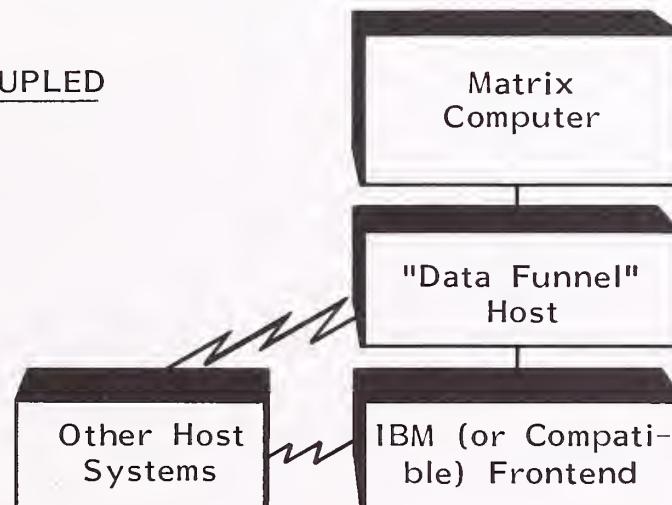
LOOSELY COUPLED OR SHARED



Amdahl 1100, 1200

IBM 308X, Amdahl 58XX
NAS 9000

VERY LOOSELY COUPLED OR NETWORKED



GRC SC 533

VAX, Perkin-Elmer, Etc.

IBM 308X, Etc.

computers, the heavy processing requirements associated with large data bases (and complex data structures) have been virtually ignored by computer architects. Also, user data base managers do not clearly understand the additional data base management burden that INPUT foresees, even though they are confronted with massive mainframe upgrades with each "advance" along the road to central, corporate data bases.

- Nevertheless, the need for data base machines to support relatively new data base structures--specifically the relational model--has been recognized as soon as relational data base systems got out of the laboratory and university and into the commercial environment. Britton Lee's Intelligent Database Machine (IDM) was introduced in 1981 specifically to support a full-blown relational DBMS. (To quote David Britton: "We had to ship the hardware that would make it perform.") The IDM operates as a "global data base resource to mainframes, minis, and microcomputers as standalones, in clusters, or on a local area network (LAN)."
- The concept of associative memories has been kicked around in the research laboratories for 20 years, and has finally begun to emerge in the attempts to solve the data base performance problem. Associative memories fall under the same general category as vector processors, being single-instruction, multiple-data machines. They can give a substantial performance boost in the data base environment; but it is not clear whether they will be isolated in separate data base machines or used in advanced storage controllers. INPUT predicts that IBM will find it necessary to improve performance of DB2 (even on Sierra), and an enhanced controller is a good bet.
- INPUT has also predicted that a "sort box" would be developed to improve DBMS performance. This particular prediction goes back seven years, when IMS was just beginning to hit the large mainframes. IBM has never done anything along this line, because ISM performance has been extremely effective in forcing mainframe upgrades. However, at the National Computer Conference in Las Vegas (July 1984), the Japanese fifth-generation project prompted two announcements:

- The sequential inference machine (SIM), which was described as a parallel-multiprocessing machine, was announced. (This system should not be confused with the large-scale parallel processors described above. It is a prototype personal computer with the power of a DECsystem-20 Model 2060.)
- A relational data base machine called Delta was announced in support of SIM. It was described as a "special engine for merging and sorting data."
- The differentiation of function for the management of data will become increasingly necessary as data bases grow in size. The development of these data base engines will unquestionably take many forms, but they will closely parallel the sample processor configurations for heavy computation. (See Exhibit II-3.) The large mainframe will be left with the role of traffic cop on the network and will be handling the flow of data and information among intelligent terminals (users), processors, and data bases.

4. NETWORK MANAGEMENT

- This remaining function looks more like an electronic switching system than a general purpose mainframe. Network management must:
 - Be highly reliable.
 - Provide routings, "busy signals," messages on use, etc.
 - Maintain directories of subscribers, terminals, processors, and data bases.
 - Provide billing for use.

- This is the area of conflict between AT&T and IBM: both must compete against the other's past strengths. However, we are already beginning to see integration of communications functions (voice, data, and image), and there is agreement that large mainframes operating under MVS/XA were not originally designed for the interactive environment. On the other hand, UNIX was not designed to handle the management of large data bases or complex scheduling problems.
- Both new hardware and new software are going to be required before the proper management of networks, including the currently confused LANs, is possible. However, for our purposes the main point is that all services will not be funneled through the bottleneck of the large mainframes.

E. MECHANIZATION AND MICROS

- Just as major functions (computation, data base management, network management) will tend to become differentiated into more specialized hardware/software systems, the mere routine functions of general purpose mainframes will tend to become mechanized.
- Language interpretations, compilation, and error checking will be performed at the intelligent workstations before moving onto the network. Languages will continue to proliferate, but terminals will tend to become more specialized, with many translators, interpreters, and compilers built into the hardware. For example, a scientist dealing with matrix-matrix multiplication will have an appropriate algebraic language built into the terminal, and someone using the query facility for DB2 will also have a specialized terminal. Neither would require intermediate translation prior to transmission to the matrix computer or relational DBMS.

- As knowledge-based systems develop, the integration of hardware, software, and knowledge base will result in a fully mechanized (single-function) system for the use of the operator. For example, a doctor will have a system for diagnosis at the office; the system will be for that purpose alone. She or he will keep track of the stock portfolio on a personal computer.
- It all adds up to the eventual demise of the large mainframe, but the transition will be relatively slow. In fact, Sierra has now been slipped to 1985, and Summit is still in the works for the late 1980s. But starting in the 1990s there will be more concern for residual values of networks rather than discrete, large-scale systems.

III RESIDUAL VALUE FORECASTS

A. FACTORS AFFECTING RESIDUAL VALUE FORECASTS

- Computer equipment residual value forecasts are based upon:
 - Analysis of historical events and trends leading to judgments about whether (and in what way) such trends may change.
 - Predictions by computer industry experts on expected actions by IBM and responding strategies by both software-compatible mainframe manufacturers and vendors of alternate terminal solutions.
 - Analysis of variables affecting residual values, as listed in Exhibit III-1.
- The most visible factor affecting IBM mainframe residual values is the announcement of a new series of large-scale systems. Computer industry experts make a living with such predictions, but actual announcements have never caused violent fluctuations in INPUT's projected residual values.
- When IBM announced the 308XX in early 1984, INPUT issued an Executive Bulletin on the probable impact on both residual values and the announcement of Sierra (these predictions were reviewed in Large-Scale Systems Directions: Midyear Update and concluded that IBM would not be under any pressure to announce a new series despite the experts' predictions for announcement late this year. The current situation is as follows:

EXHIBIT III-1

FACTORS AFFECTING COMPUTER EQUIPMENT RESIDUAL VALUES

- IBM practices and policies
 - New product announcements
 - . Price/performance ratios relative to existing products.
 - . Ease of conversions, transitions, and lead time in obtaining new products.
 - . Ease of installation and maintenance.
 - . Effect on perceptions about IBM's technical direction.
 - Pricing policies
 - . Price increases or decreases on existing products.
 - . Rental versus purchase break-even ratios.
 - . Lease plans and penalty provisions for lease termination.
 - . Purchase option accruals.
 - Maintenance policies
 - . Availability and cost.
 - . Attitude toward other vendor modifications to IBM equipment.
- Alternative equipment services
 - Price/performance of plug- (software-) compatible alternatives.
 - Third-party leasing options.
- Other variables
 - Environmental support considerations, e.g., electrical power consumption, air conditioning needs, space requirements.
 - Tax considerations, e.g., income tax incentives such as investment tax credit and accelerated depreciation, and also property taxation rates.
 - General economic conditions, e.g., cost and availability of capital and overall demand for computing capacity.

- As the year has progressed, the "experts" have begun to push back their predictions. One such expert was recently quoted as stating that Sierra would be announced "perhaps as early (?) as third-quarter 1985," with Summit being announced by the second quarter of 1987.
 - . When predictions are changed, it is always good strategy to pretend you were not one of the experts who had been predicting an earlier announcement. "Perhaps as early" is a good example of a not too subtle attempt at disassociating from the rest of the pack.
 - . INPUT titled its Executive Bulletin on the 308XX models Musical Mainframes Again . . ., but positioning Sierra and Summit less than two years apart would seem to indicate that IBM has been leaking "information" to the experts and that "orchestrated confusion" really reigns.
 - . It is obvious INPUT cannot let predictions by outside experts play too important a role in our forecasts of residual values, even though we do read them to complement the more serious analysis.
- One of the sources we use is a poll of what IBM is telling its major customers concerning future announcements. Although this method does not contribute to significantly more accurate predictions than listening to the experts, it is informative to follow the explanations used when IBM shifts direction. For example, as recently as three or four months ago, IBM was still telling some major customers that Sierra would be announced in the fourth quarter of 1984. Now, IBM is leaking information to its customers about "terminal problems" with Sierra and encouraging the purchase of upgrades to the 3084Q.

- Any new system or technology can be reported to have technical problems if the manufacturer doesn't want to announce the system. It is an extremely convenient excuse for IBM marketing representatives because they cannot be expected either to understand or to explain the supposed technical problem.
 - Technical problems seldom delay announcements when competitive pressures warrant getting a product in the marketplace. (Deliveries may be delayed, but announcements of new products will not be delayed.)
 - Therefore, INPUT does not rely too much on what IBM is telling its customers.
- The truth of the matter is, when anticipating IBM announcements, INPUT relies primarily upon what makes good financial sense for IBM. That is the secret of INPUT's early 1984 prediction: "Announcement and/or delivery schedules of Sierra will be delayed to permit the 308XX profitability advantages to be exploited." It is still INPUT's opinion that this is a more plausible reason than "technical problems" for the Sierra slippage. The slippage has been factored into our residual value forecasts.

B. ANNOUNCEMENTS

- The following announcements of significance in forecasting mainframe residual values have been made since the Large-Scale Systems Directions: Midyear Update was published.
 - IBM announced price reductions on the 3080 series of mainframes in early September of 1984.

- The price cuts ranged from 12-16% on processors.
 - Memory prices were reduced from \$20,000 to \$16,250 per mega-byte, with 8 M-bytes being reduced from \$160,000 to \$130,000.
 - The prices of 3082 processor controllers were also reduced, with the low-end Model X8 going down to \$145,000 from \$170,000 and the high-end Model X48 going to \$490,000 from \$540,000.
 - The 3087 Model I coolant distribution unit was cut from \$60,000 to \$50,000.
 - And the price of I/O channels was cut from \$18,750 to \$16,250.
- Within a week, both Amdahl and NAS reacted.
- Amdahl reduced purchase prices on all 580 series mainframes from 9% to 16%.
 - NAS announced reduced purchase prices on their 8000/9000/9100 series from 12% to 16%.
- While both NAS and Amdahl were announcing the supercomputers that were reviewed in the preceding section of this report, IBM was keeping the pot boiling at the low end of the 4300 line and at the crucial overlap between the high end of the 4300 and 3080 series.
- In mid-September, IBM announced a low-end addition to the 4361 series, the Model 3, priced at \$56,500 with 2 M-bytes of main memory. At the same time:
- The 4321 and 4331 series of mainframes were cancelled, with sharply reduced prices (21% to 37% reductions) being in effect while orders were still being accepted (through 12/31/84).

- Existing models of the 4361 were reduced in price by 10%.
- Then in late October, IBM announced the low-end 3083 Model CX and the high-end 4381 Model Group 2.
 - A 3083 Model CX--including CPU, 3082 processor controller, 3087 Model I coolant distribution unit, 8 M-bytes of main memory, and eight channels--is priced at \$830,000.
 - A 4381 Model Group 3 processor, with 8 M-bytes of main memory and twelve standard channels, is priced at \$825,000.
 - The 3083 Model CX is reported to perform at about 2.5 mips, and the 4381 Model Group 3 at between 4.6 to 5.1 mips.
 - If you are wondering why anyone would buy a 3083 Model CX when the 4381 Model Group 3 has nearly two times the price-performance, you are in good company. Even IBM doesn't really explain.
 - However, INPUT believes the answer is clear: you buy the 3083 when you are establishing a major SNA host mode (it can grow to a 3084Q if you like), and you put in a 4381 at a mode with more predictable processing requirements (a branch office or manufacturing plant).
 - In other words, if you anticipate rapid growth you are going to pay for it--especially until Sierra is announced and large mainframes achieve price-performance more competitive with the mid-range 4300 series.

- Other announcements of general interest include:
 - IBM's acquisition of ROLM.
 - CDC getting out of the OEM peripherals business; STC filing for reorganization under Chapter II; and Honeywell selecting IBM as an OEM supplier of 3380-type disk systems. It doesn't require a lot of analysis to determine the direction of large-scale disk storage systems, and the whole sequence of events makes us more than a little melancholy. The full impact of the developments will be analyzed in the next Large-Scale Systems Directions: Disk, Tape, and Printer Systems, which will be published early in 1985.
 - Amdahl has demonstrated substantial activity in the software area recently.
 - . Improved delivery schedules on XA for 580 dual processors were announced in August (from the originally scheduled delivery time of the second quarter of 1985 to the fourth quarter of 1984).
 - . Amdahl has announced an agreement to market Ingress (a relational data base management system from Relational Technology) under UTS (Amdahl's version of UNIX for large mainframes).
 - . ASPEN (Amdahl's alternative operating system, which has been under development for several years) has not officially been announced but has been spoken of among prospective customers.
 - . And, as this report was being completed, a "multiple domain feature" hardware enhancement for the 580 series was announced, which permits two operating systems to coexist on the same uniprocessor or multiprocessor.

- INPUT, in the belief that alternatives to IBM's software strategy were highly desirable, endorsed selective deviations among plug-compatible mainframe vendors years ago. The only question INPUT has now is whether the window of opportunities has closed.

C. PROJECTED USED MARKET PRICES AND RESIDUAL VALUES

- Exhibits III-2 and III-3 contain projected used market retail values in dollars, and the projected residual values as a percent of vendor list price. It should be understood that, at any given time, three price levels exist.
 - "Retail price" is the amount an end user would pay for the equipment.
 - "Dealer price" is the amount a dealer would pay another dealer for the equipment.
 - "Wholesale price" is the amount a dealer would pay to acquire equipment for resale.
 - The dollar spread between levels is a function of the total value of the transaction. For large processors the wholesale price will typically be 80-95% of the retail price; for smaller processors 70-90% is more likely.
- Exhibits III-4 through III-24 graph the range of anticipated values (as a percent of list price) for 1986 through 1990 for the following processors:
 - IBM 4331-K02, 4341-L02, 4361-K04, 4381-K05, 3083-EX8, 3083-BX16, 3083-JX32, 3081-GX16, 3081-KX24, and 3084-QX64.
 - Amdahl 5850-24, 5860-24, 5868-32, 5870-32, and 5880-64.
 - NAS AS/6630, AS/6660, AS/8023, AS/8083, AS/9050, and AS/9070.

EXHIBIT III-2

PROJECTED USED MARKET RETAIL VALUES

VENDOR	PROCESSOR MODEL	CURRENT LIST * 11/1/84	ESTIMATED MARKET VALUE* 1985	PROJECTED USED MARKET RETAIL VALUE* AT JAN. 1 OF:				
				1986	1987	1988	1989	1990
IBM	4331-J01	60920	17058	9138	6092	4264	2437	609
	4331-K02	49205	30507	22142	13777	5905	3444	1968
	4341-K01	184500	27675	18450	14760	9225	3690	1845
	4341-L02	312000	65520	46800	34320	21840	9360	3120
	4361-0L3	71500	71500	53625	44330	34320	15730	5720
	4361-ML4	215000	172000	133300	124700	86000	23650	10750
	4381-0M2	540000	459000	437400	361800	243000	86400	43200
	4381-0M3	825000	825000	701250	618750	453750	247500	123750
	3033-N08	1274000	38220	12740	6370	3822	0	0
	3033-U12	1764000	88200	52920	26460	17640	8820	0
	3083-CX8	635000	635000	412750	254000	139700	63500	25400
	3083-EX8	810000	607500	445500	267300	121500	48600	16200
	3083-BX16	1460000	1138800	832200	525600	277400	116800	29200
	3083-JX32	2160000	1728000	1296000	864000	453600	216000	64800
	3081-GX16	2475000	2029500	1608750	990000	495000	297000	123750
	3081-KX24	3365000	2860250	2288200	1514250	807600	504750	235550
	3084-QX64	6010000	5409000	4507500	3606000	2704500	1502500	721200
AMDAHL	470-V7	By Quote	-	-	-	-	-	-
	470-V8	By Quote	-	-	-	-	-	-
	5840-16	1700000	1360000	1020000	646000	289000	102000	17000
	5850-24	2140000	1712000	1284000	834600	385200	171200	21400
	5860-24	2560000	2176000	1740800	1075200	512000	256000	76800
	5867-32	3360000	3024000	2352000	1478400	739200	369600	100800
	5868-32	3690000	3321000	2583000	1623600	811800	405900	110700
	5870-32	3930000	3340500	2751000	1886400	1021800	510900	235800
	5880-64	5040000	4284000	3528000	2520000	1411200	907200	453600

*Dollars

EXHIBIT III-2 (Cont.)

PROJECTED USED MARKET RETAIL VALUES

VENDOR	PROCESSOR MODEL	CURRENT LIST* 11/1/84	ESTIMATED MARKET VALUE* 1985	PROJECTED USED MARKET RETAIL VALUE*AT JAN. 1 OF:				
				1986	1987	1988	1989	1990
NAS†	AS/6620	255000	153000	76500	35700	20400	12750	2550
	AS/6630	341500	211730	119525	61470	27320	17075	6830
	AS/6650	417500	279725	167000	104375	62525	41750	16700
	AS/6660	475000	475000	403750	346750	213750	85500	47500
	AS/8023-8	699000	594150	363480	230670	83880	41940	13980
	AS/8043-8	1067000	853600	586850	373450	160050	85360	42680
	AS/8053-8	1492000	1193600	850440	537120	223800	134280	74600
	AS/8063-8	1905000	1714500	1143000	742950	323850	209550	133350
	AS/8083-16	3074000	2705120	2151800	1321820	614800	430360	245920
	AS/9040-8	1492000	596800	298400	134280	74600	29840	14920
	AS/9050-8	1909000	801780	477250	229080	133630	76360	38180
	AS/9060-16	2308000	1038600	692400	415440	230800	138480	92320
	AS/9070-16	3249000	1624500	1364580	779760	487350	259920	194940
	AS/9080-16	4140000	2152800	1945800	1242000	745200	414000	289800

*Dollars

†National Advanced Systems (NAS) does not quote processor prices separately; list price on this schedule includes power distribution unit, controller and console, where appropriate.

EXHIBIT III-3

**PROJECTED RESIDUAL VALUES
(AS A PERCENT OF VENDOR LIST PRICE)**

VENDOR	PROCESSOR MODEL	CURRENT LIST* 11/1/84	ESTIMATED MARKET VALUE 1985	PROJECTED RESIDUAL VALUE AS A PERCENT OF VENDOR LIST PRICE VALUE AT JAN. 1 OF:				
				1986	1987	1988	1989	1990
IBM	4331-J01	60920	28%	15%	10%	7%	4%	1%
	4331-K02	49205	62%	45%	28%	12%	7%	4%
	4341-K01	184500	15%	10%	8%	5%	2%	1%
	4341-L02	312000	21%	15%	11%	7%	3%	1%
	4361-DL3	71500	100%	75%	62%	48%	22%	8%
	4361-ML4	230000	80%	62%	58%	40%	11%	5%
	4381-DM2	540000	85%	81%	67%	45%	16%	8%
	4381-DM3	825000	100%	85%	75%	55%	30%	15%
	3033-N08	1274000	3%	1%	1%	0%	0%	0%
	3033-U12	1764000	5%	3%	2%	1%	1%	0%
	3083-CX8	635000	100%	65%	40%	22%	10%	4%
	3083-EX8	810000	75%	55%	33%	15%	6%	2%
	3083-BX16	1460000	78%	57%	36%	19%	8%	2%
	3083-JX32	2160000	80%	60%	40%	21%	10%	3%
	3081-GX16	2475000	82%	65%	40%	20%	12%	5%
	3081-KX24	3365000	85%	68%	45%	24%	15%	7%
	3084-QX64	6010000	90%	75%	60%	45%	25%	12%
AMDAHL	470-V7	By Quote	2%	1%	1%	0%	0%	0%
	470-V8	By Quote	4%	2%	1%	1%	1%	0%
	5840-16	1700000	80%	60%	38%	17%	6%	1%
	5850-24	2140000	80%	60%	39%	18%	8%	1%
	5860-24	2560000	85%	66%	42%	20%	10%	3%
	5867-32	3360000	90%	70%	44%	22%	11%	3%
	5868-32	3690000	90%	70%	44%	22%	11%	3%
	5870-32	3930000	85%	70%	48%	26%	13%	6%
	5880-64	5040000	85%	70%	50%	28%	18%	9%

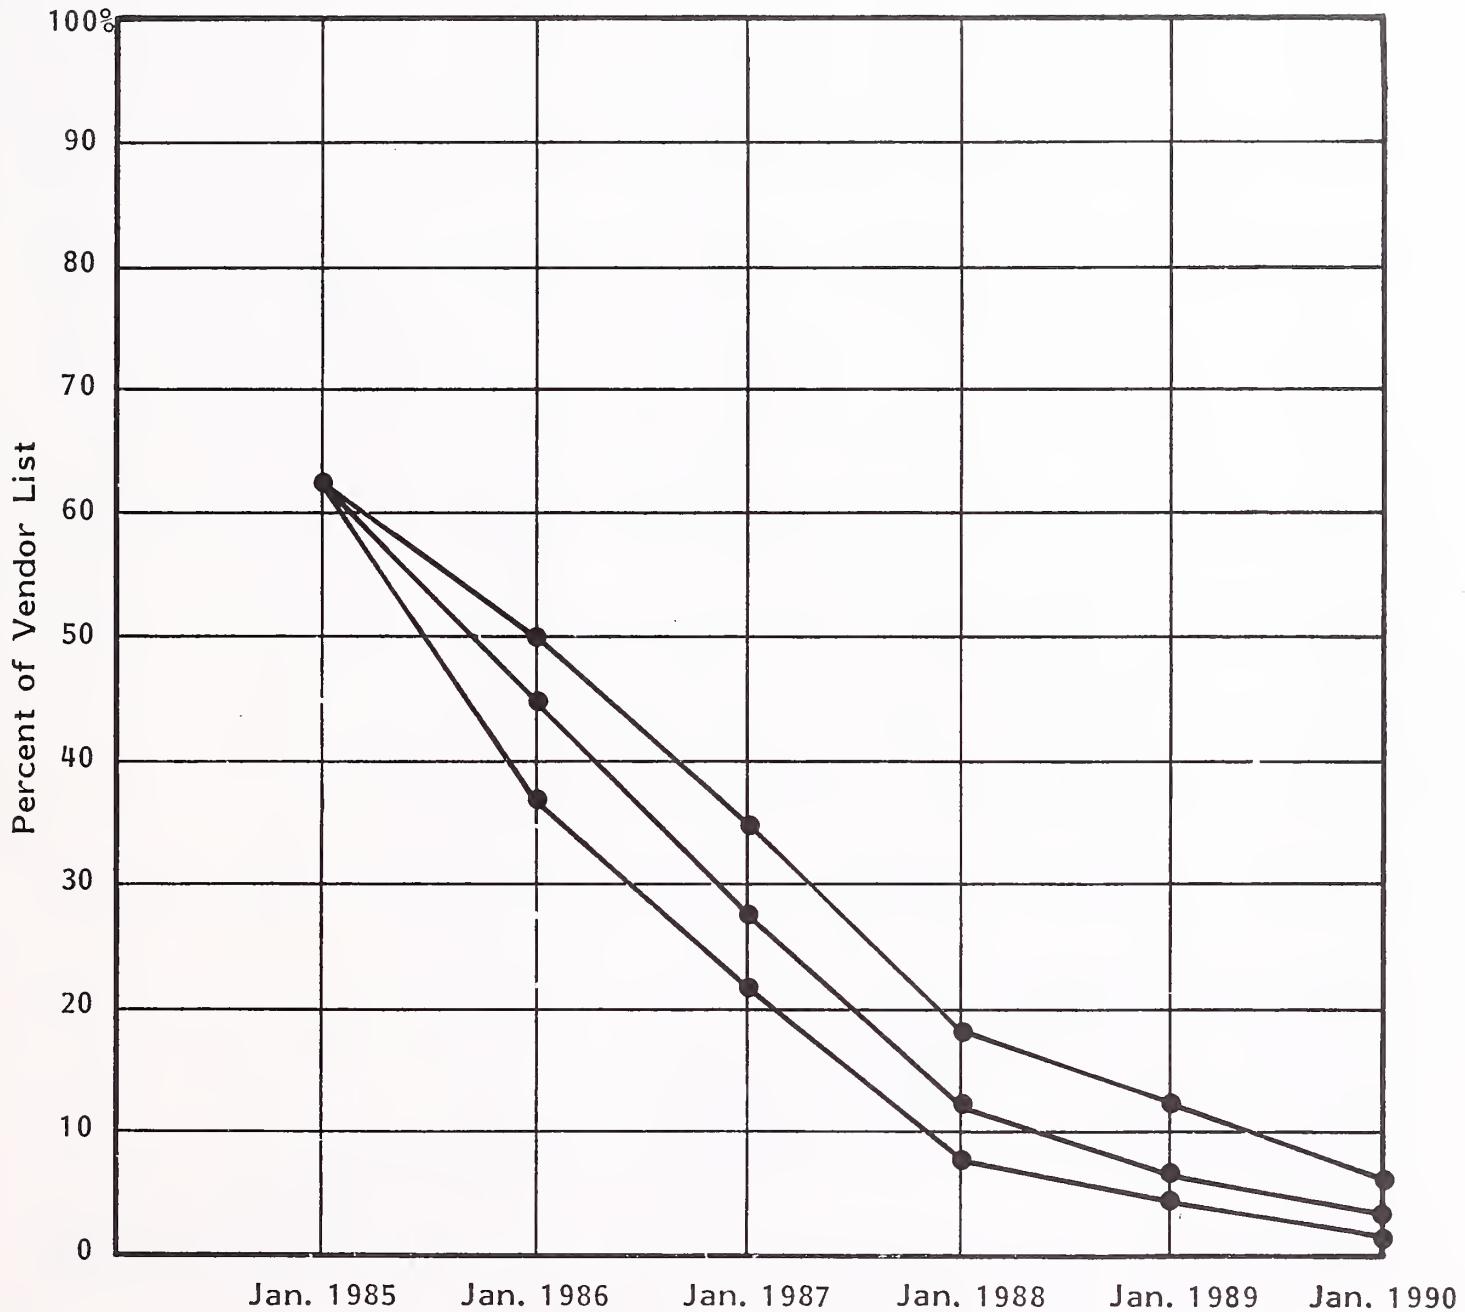
EXHIBIT III-3 (Cont.)

PROJECTED RESIDUAL VALUES
(AS A PERCENT OF VENDOR LIST PRICE)

VENDOR	PROCESSOR MODEL	CURRENT LIST 11/1/84	ESTIMATED MARKET VALUE 1985	PROJECTED RESIDUAL VALUE AS A PERCENT OF VENDOR LIST PRICE VALUE AT JAN. 1 OF:				
				1986	1987	1988	1989	1990
NAS*	AS/6620	255000	60%	30%	14%	8%	5%	1%
	AS/6630	341500	62%	35%	18%	8%	5%	2%
	AS/6650	417500	67%	40%	25%	15%	10%	4%
	AS/6660	475000	100%	60%	35%	22%	15%	7%
	AS/8023-8	699000	85%	52%	33%	12%	6%	2%
	AS/8043-8	1067000	80%	55%	35%	15%	8%	4%
	AS/8053-8	1492000	80%	57%	36%	15%	9%	5%
	AS/8063-8	1905000	90%	60%	39%	17%	11%	7%
	AS/8083-16	3074000	88%	70%	43%	20%	14%	8%
	AS/9040-8	1492000	40%	20%	9%	5%	2%	1%
	AS/9050-8	1969000	42%	25%	12%	7%	4%	2%
	AS/9060-16	2308000	45%	30%	18%	10%	6%	4%
	AS/9070-16	3249000	50%	42%	24%	15%	8%	6%
	AS/9080-16	4140000	52%	47%	30%	18%	10%	7%

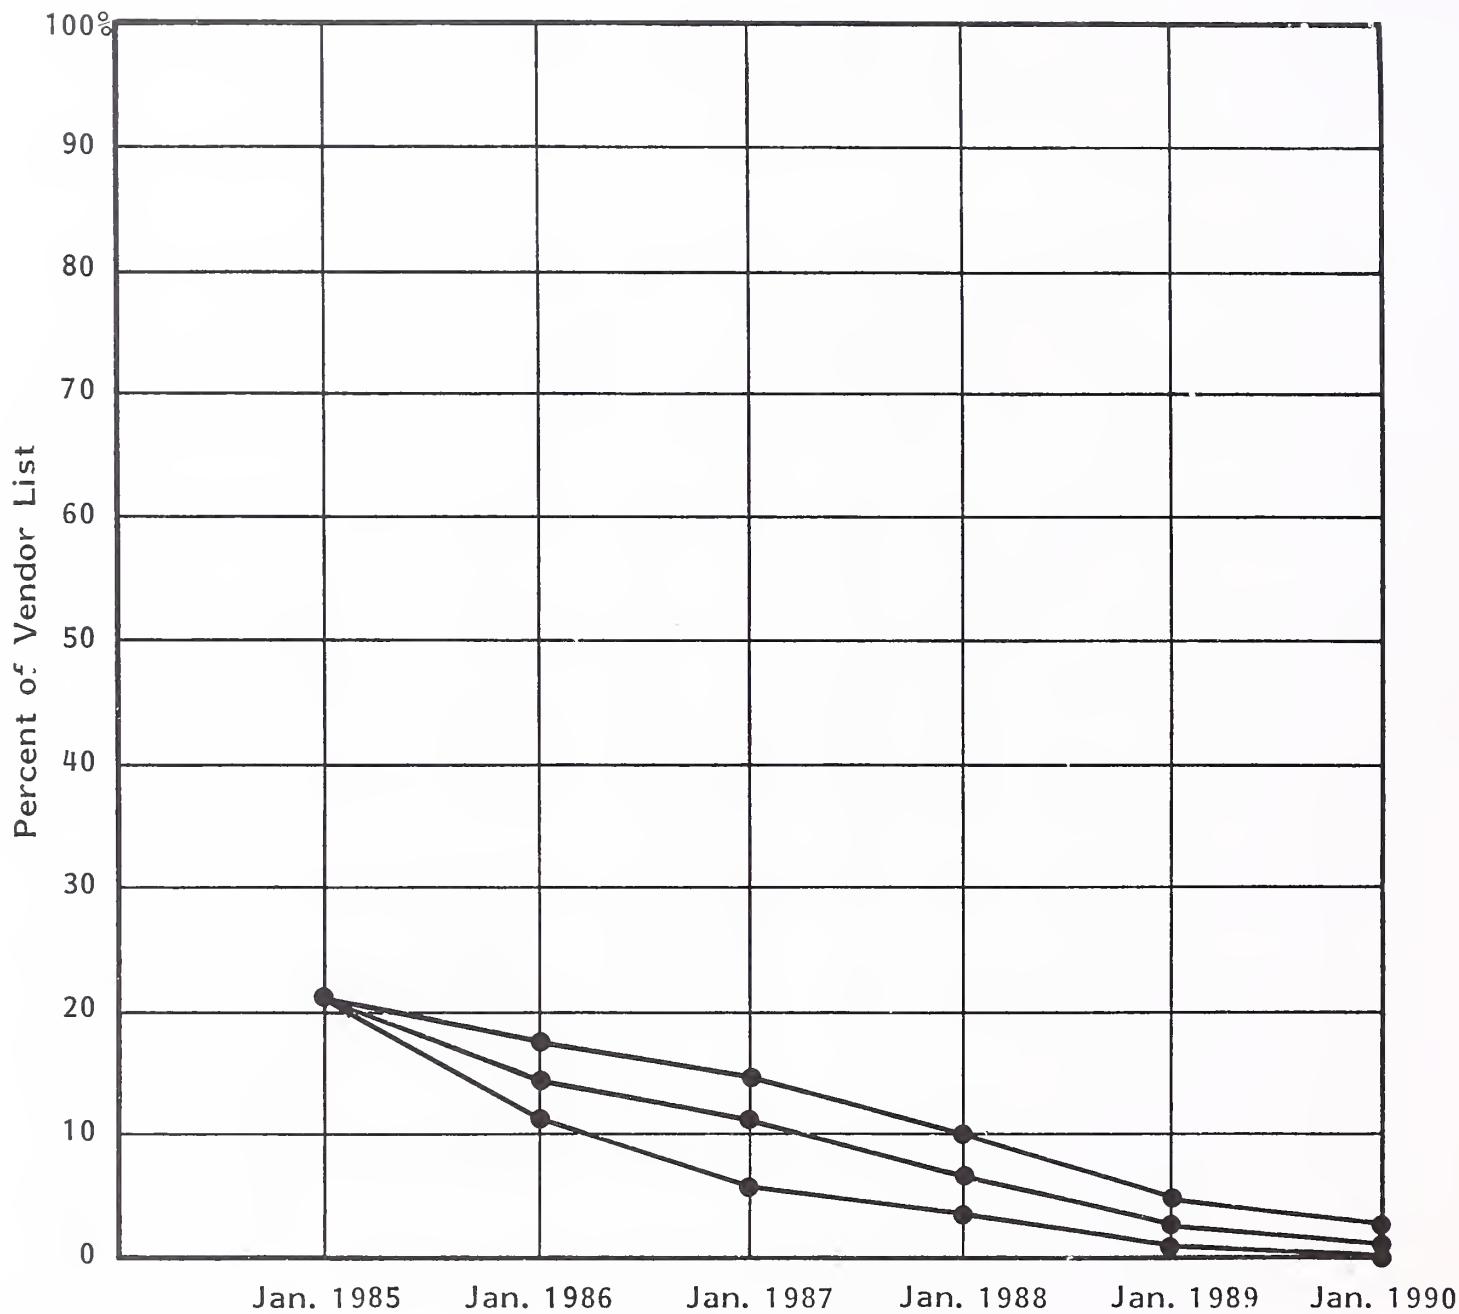
*National Advanced Systems (NAS) does not quote processor prices separately; list price on this schedule includes power distribution unit, controller and console, where appropriate.

EXHIBIT III-4
RESIDUAL VALUE FORECAST FOR
IBM 4331-K02 PROCESSOR



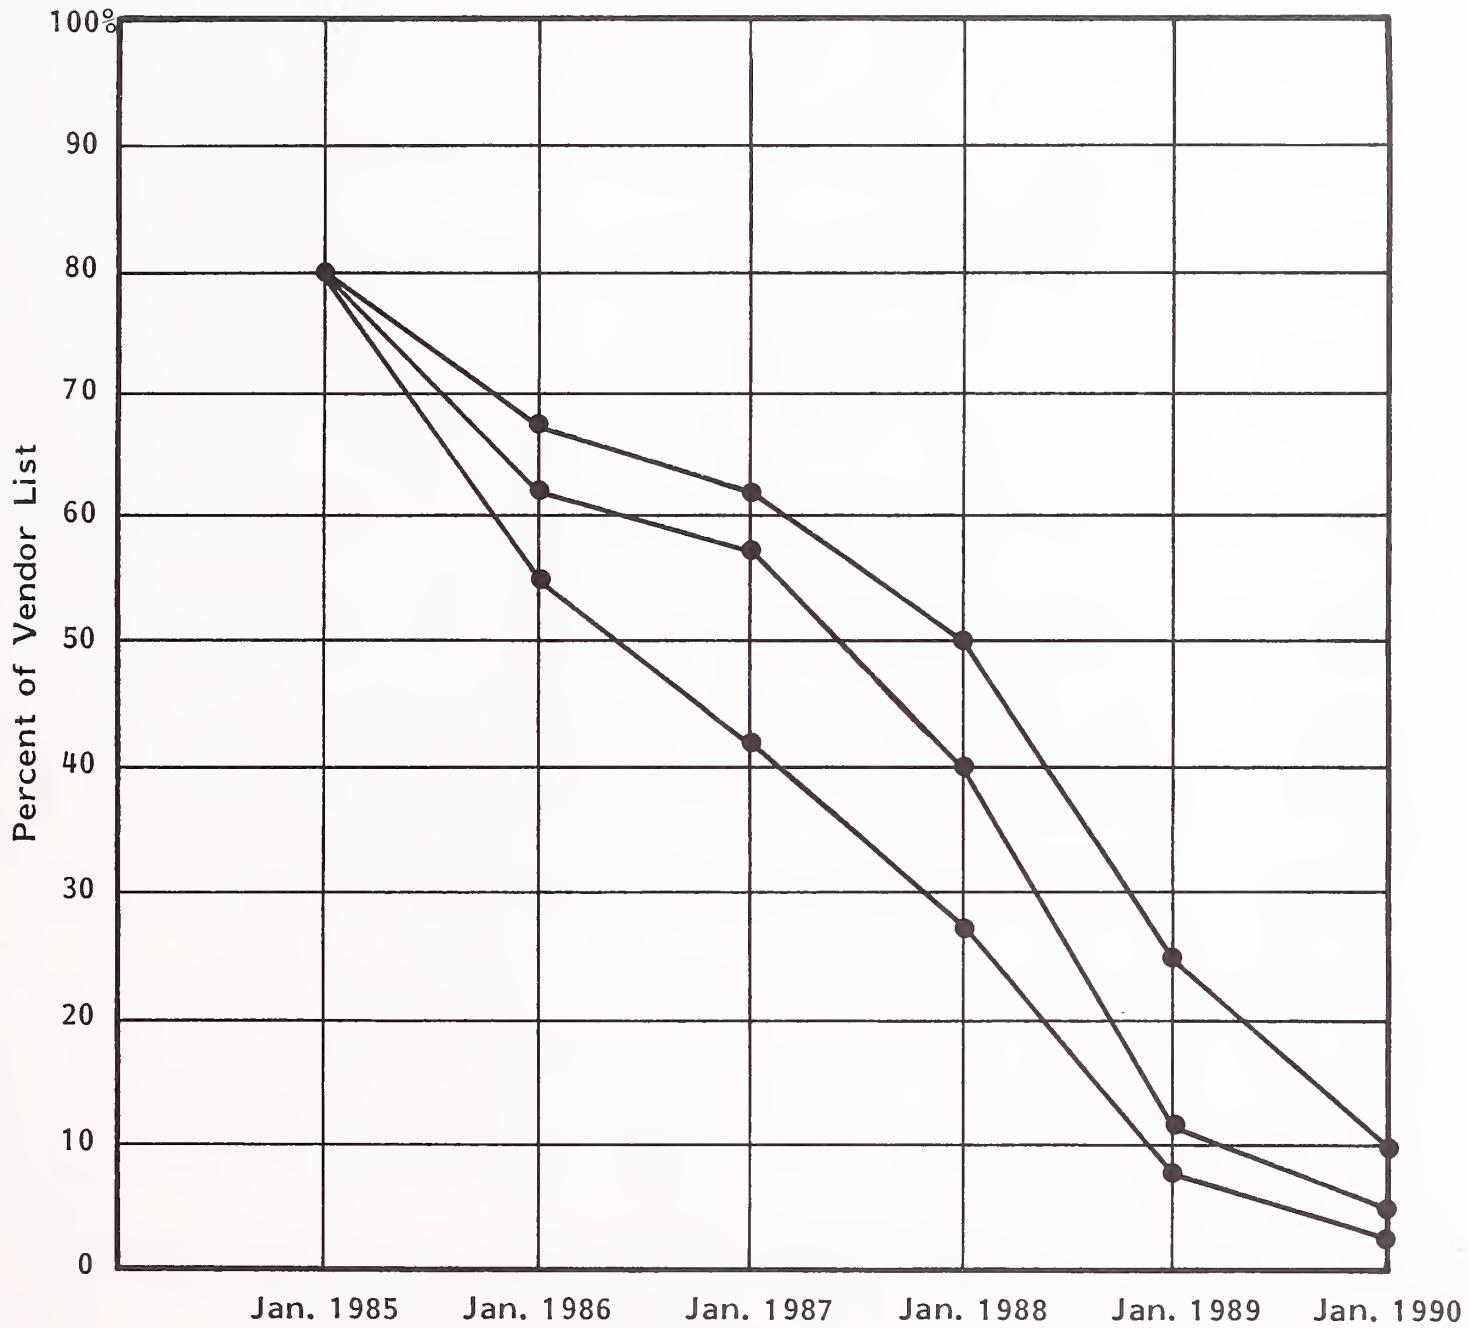
PROJECTED VALUES RANGE	JAN. 1986	JAN. 1987	JAN. 1988	JAN. 1989	JAN. 1990
High	50%	35%	18%	12%	7%
Expected	45	28	12	7	4
Low	37	22	8	5	2

EXHIBIT III-5
 RESIDUAL VALUE FORECAST FOR
 IBM 4341-L02 PROCESSOR



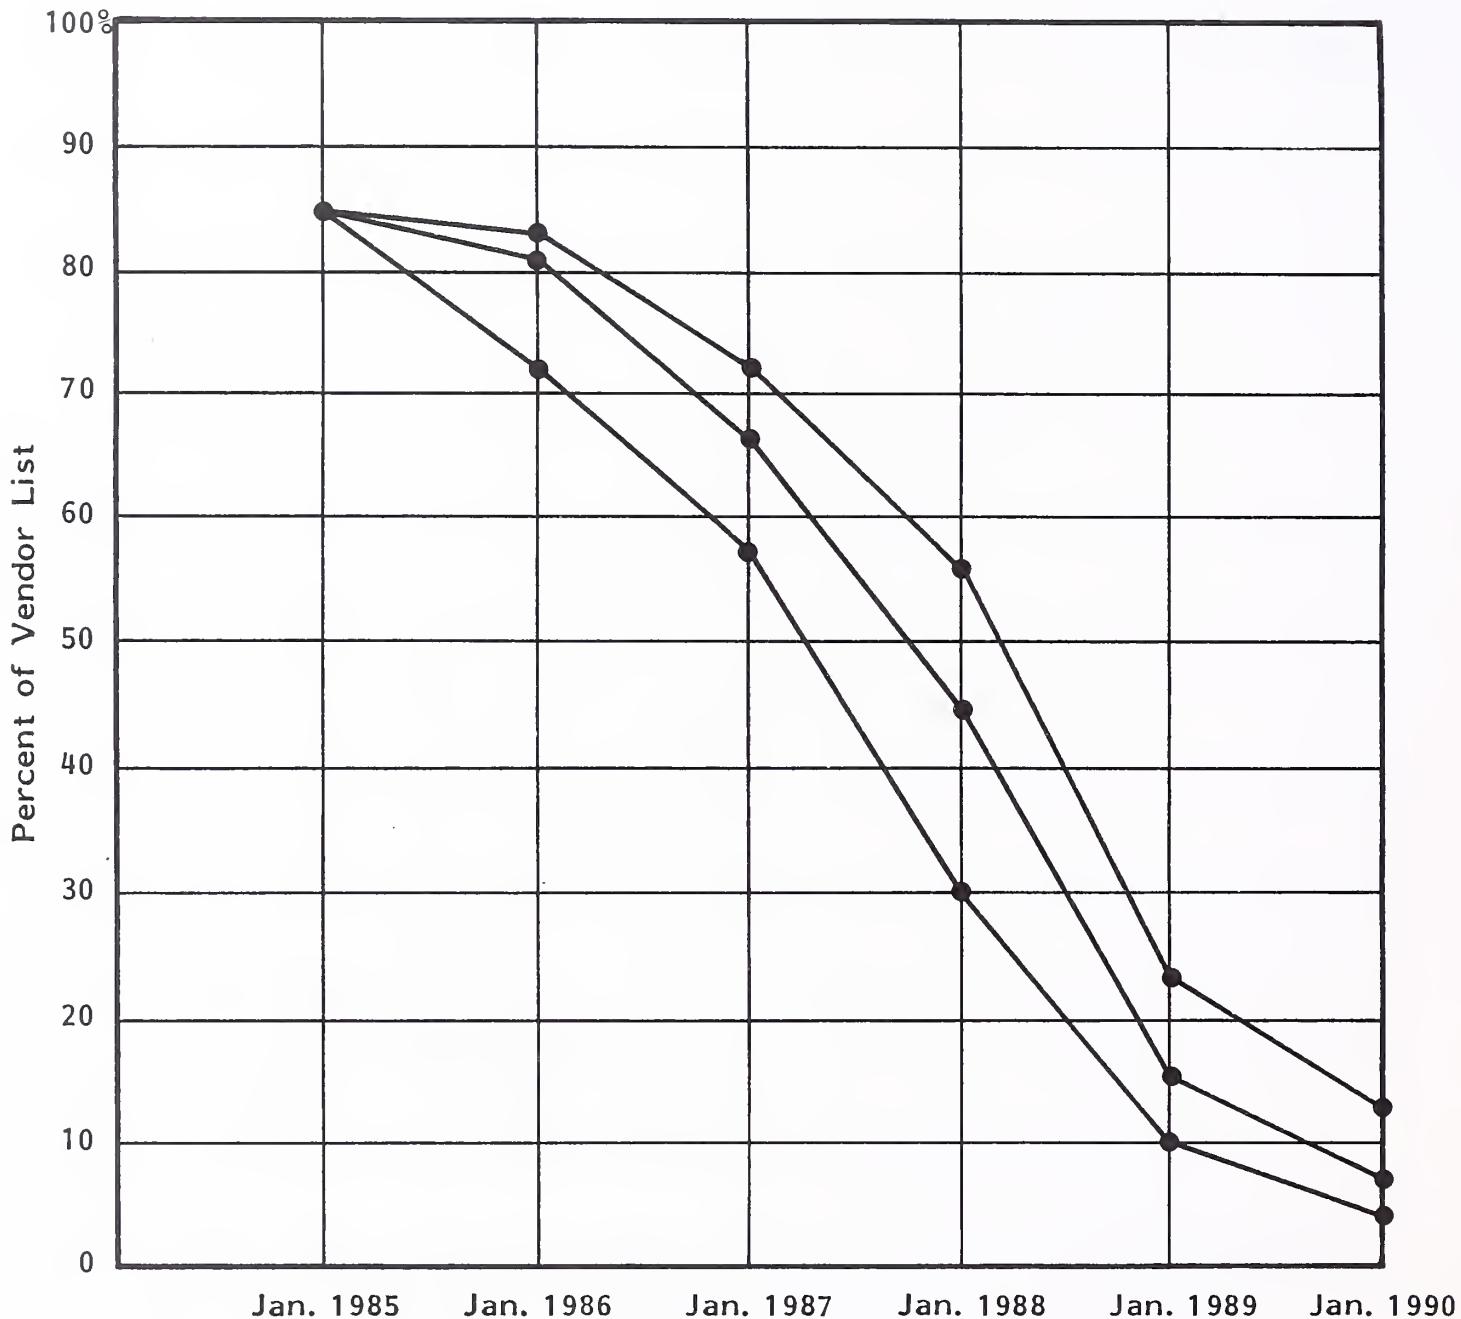
PROJECTED VALUES RANGE	JAN. 1986	JAN. 1987	JAN. 1988	JAN. 1989	JAN. 1990
High	18%	10%	10%	5%	3%
Expected	15	11	7	3	1
Low	11	6	4	1	0

EXHIBIT III-6
RESIDUAL VALUE FORECAST FOR
IBM 4361-KO4 PROCESSOR



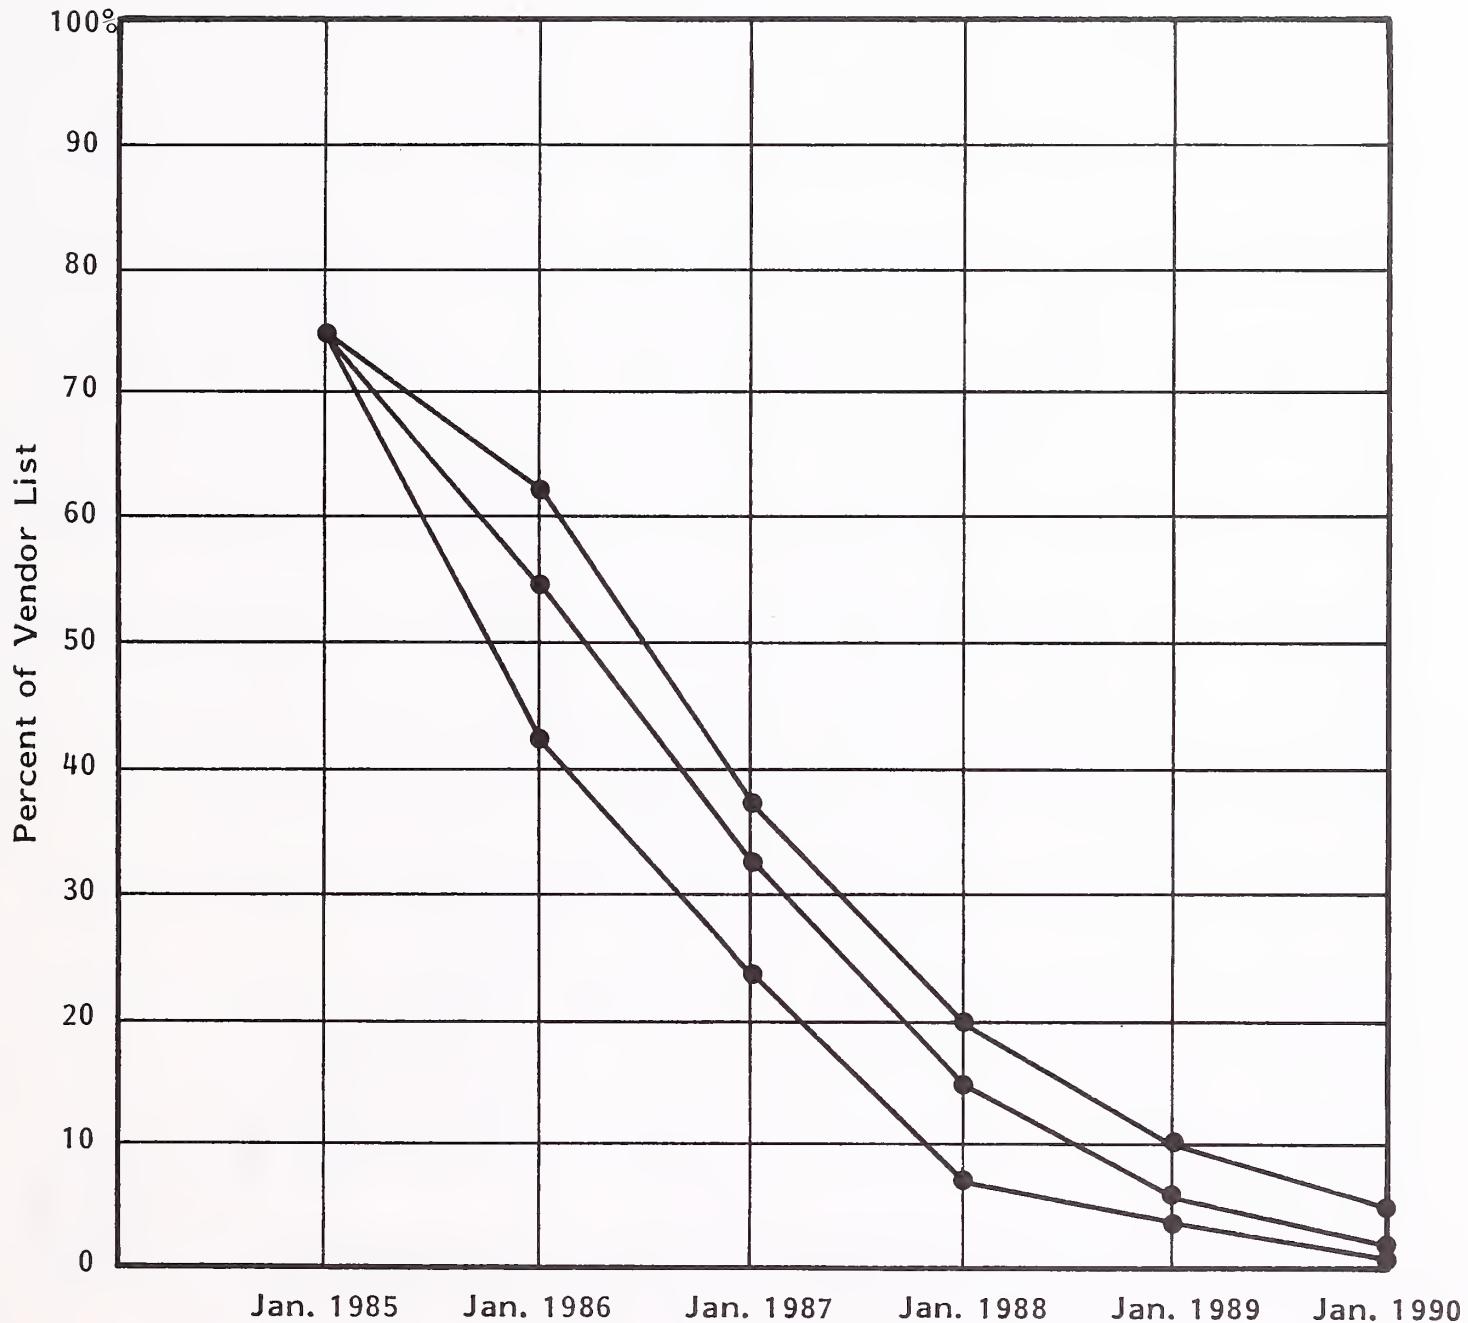
PROJECTED VALUES RANGE	JAN. 1986	JAN. 1987	JAN. 1988	JAN. 1989	JAN. 1990
High	68%	62%	50%	25%	10%
Expected	62	58	40	11	5
Low	55	42	28	8	3

EXHIBIT III-7
RESIDUAL VALUE FORECAST FOR
IBM 4381-KO5 PROCESSOR



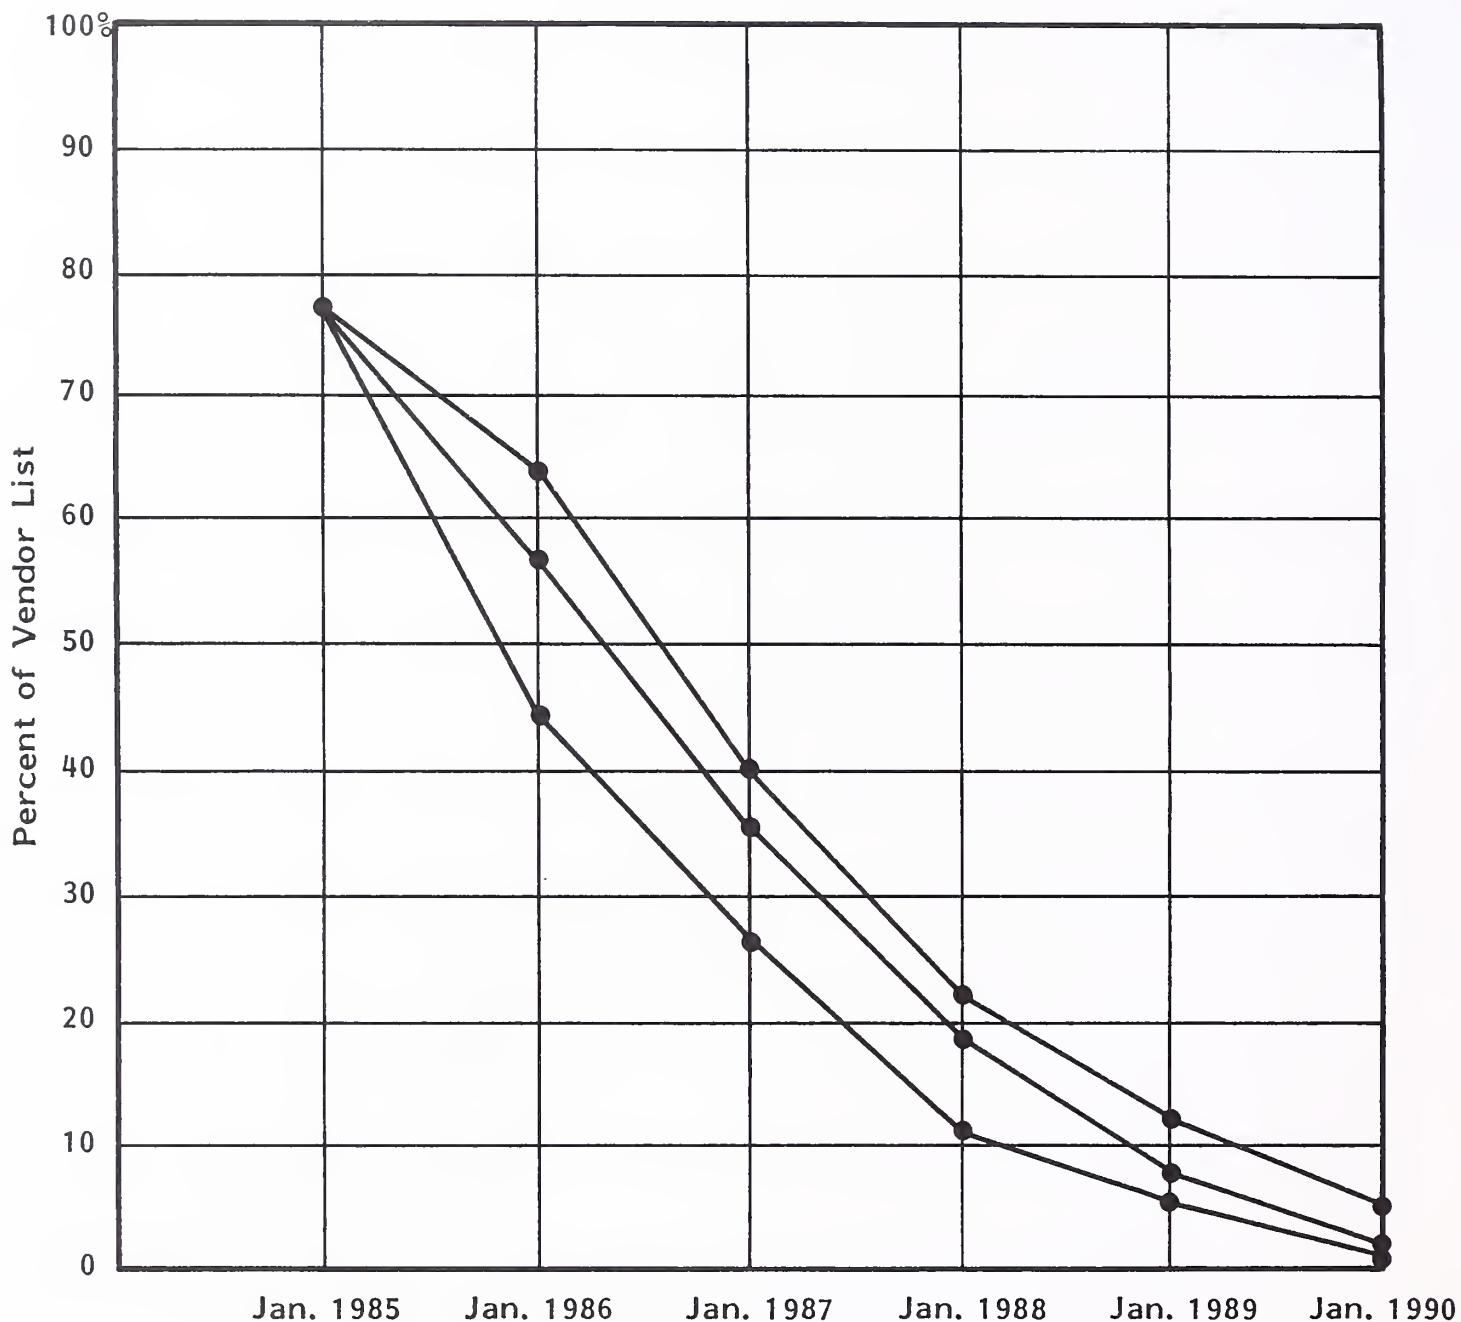
PROJECTED VALUES RANGE	JAN. 1986	JAN. 1987	JAN. 1988	JAN. 1989	JAN. 1990
High	83%	72%	57%	23%	12%
Expected	81	67	45	16	8
Low	72	58	30	10	5

EXHIBIT III-8
RESIDUAL VALUE FORECAST FOR
IBM 3083-EX8 PROCESSOR



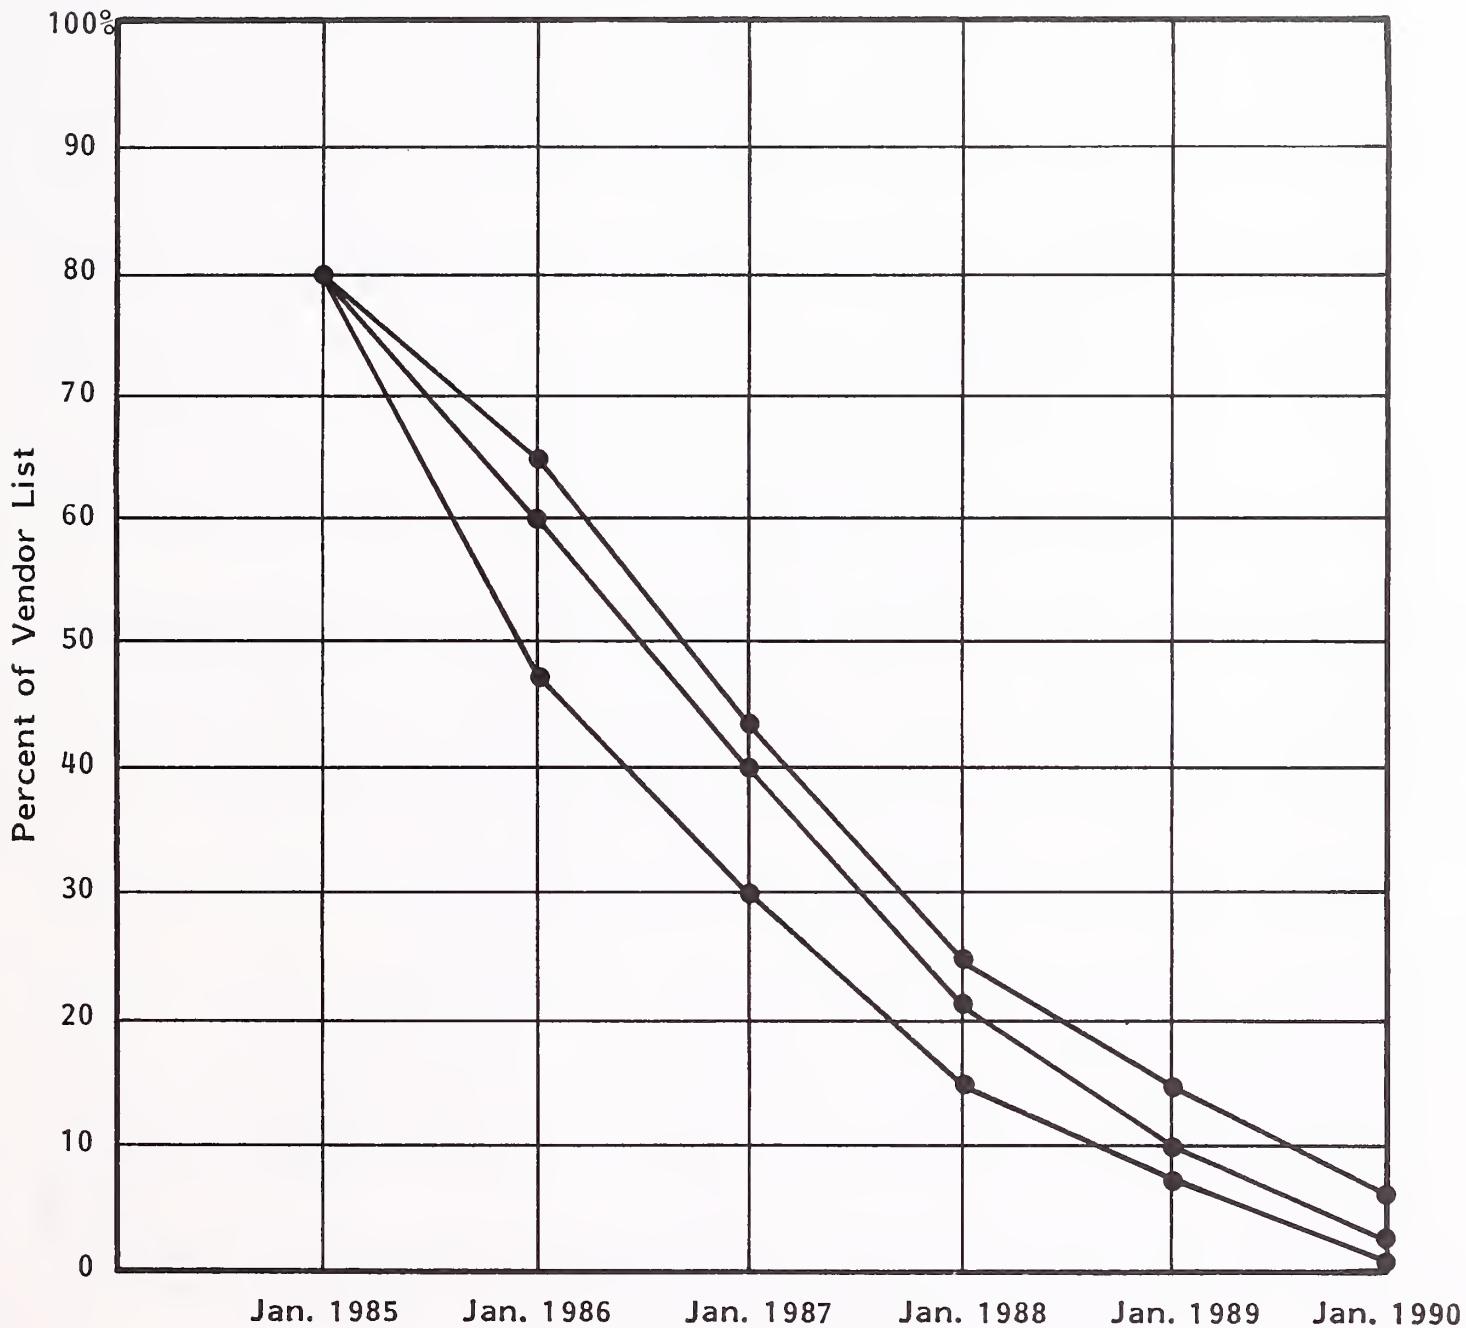
PROJECTED VALUES RANGE	JAN. 1986	JAN. 1987	JAN. 1988	JAN. 1989	JAN. 1990
High	62%	38%	20%	10%	5%
Expected	55	33	15	6	2
Low	42	24	8	4	1

EXHIBIT III-9
RESIDUAL VALUE FORECAST FOR
IBM 3083-BX16 PROCESSOR



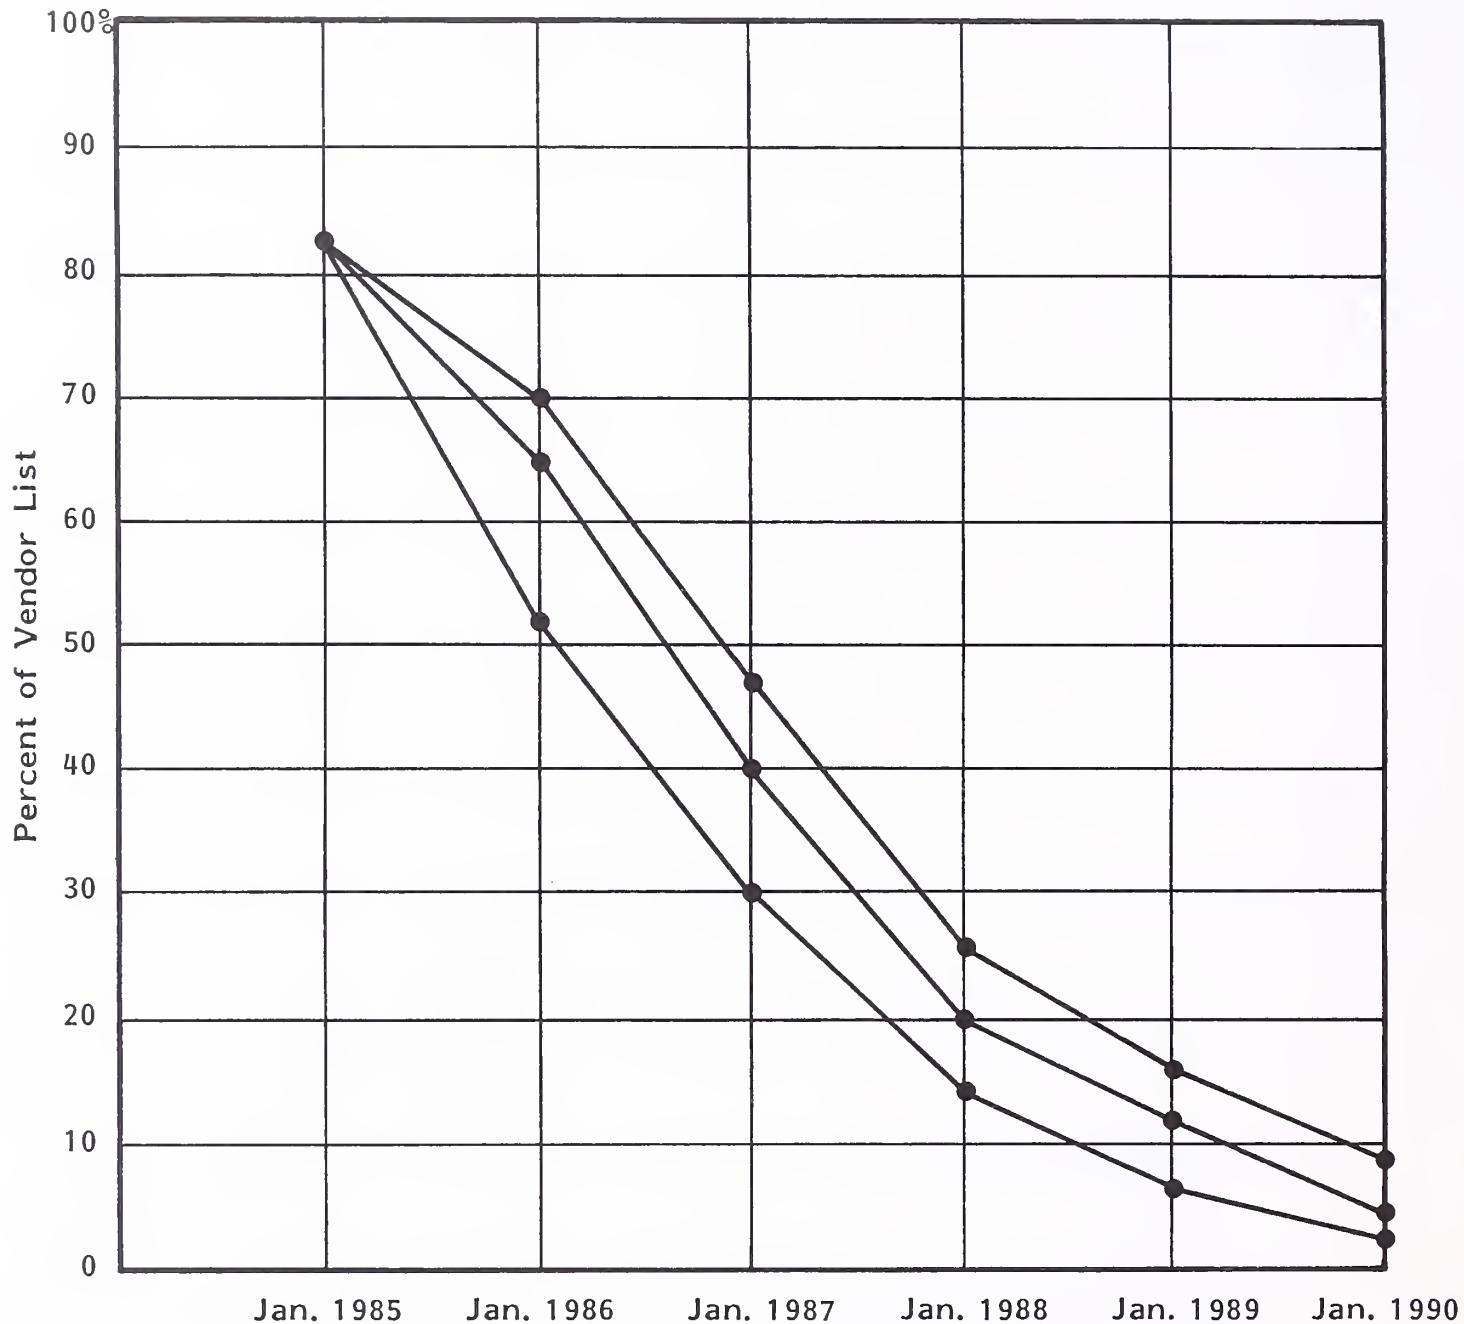
PROJECTED VALUES RANGE	JAN. 1986	JAN. 1987	JAN. 1988	JAN. 1989	JAN. 1990
High	64%	40%	22%	12%	6%
Expected	57	36	19	8	2
Low	45	27	11	6	1

EXHIBIT III-10
RESIDUAL VALUE FORECAST FOR
IBM 3083-JX32 PROCESSOR



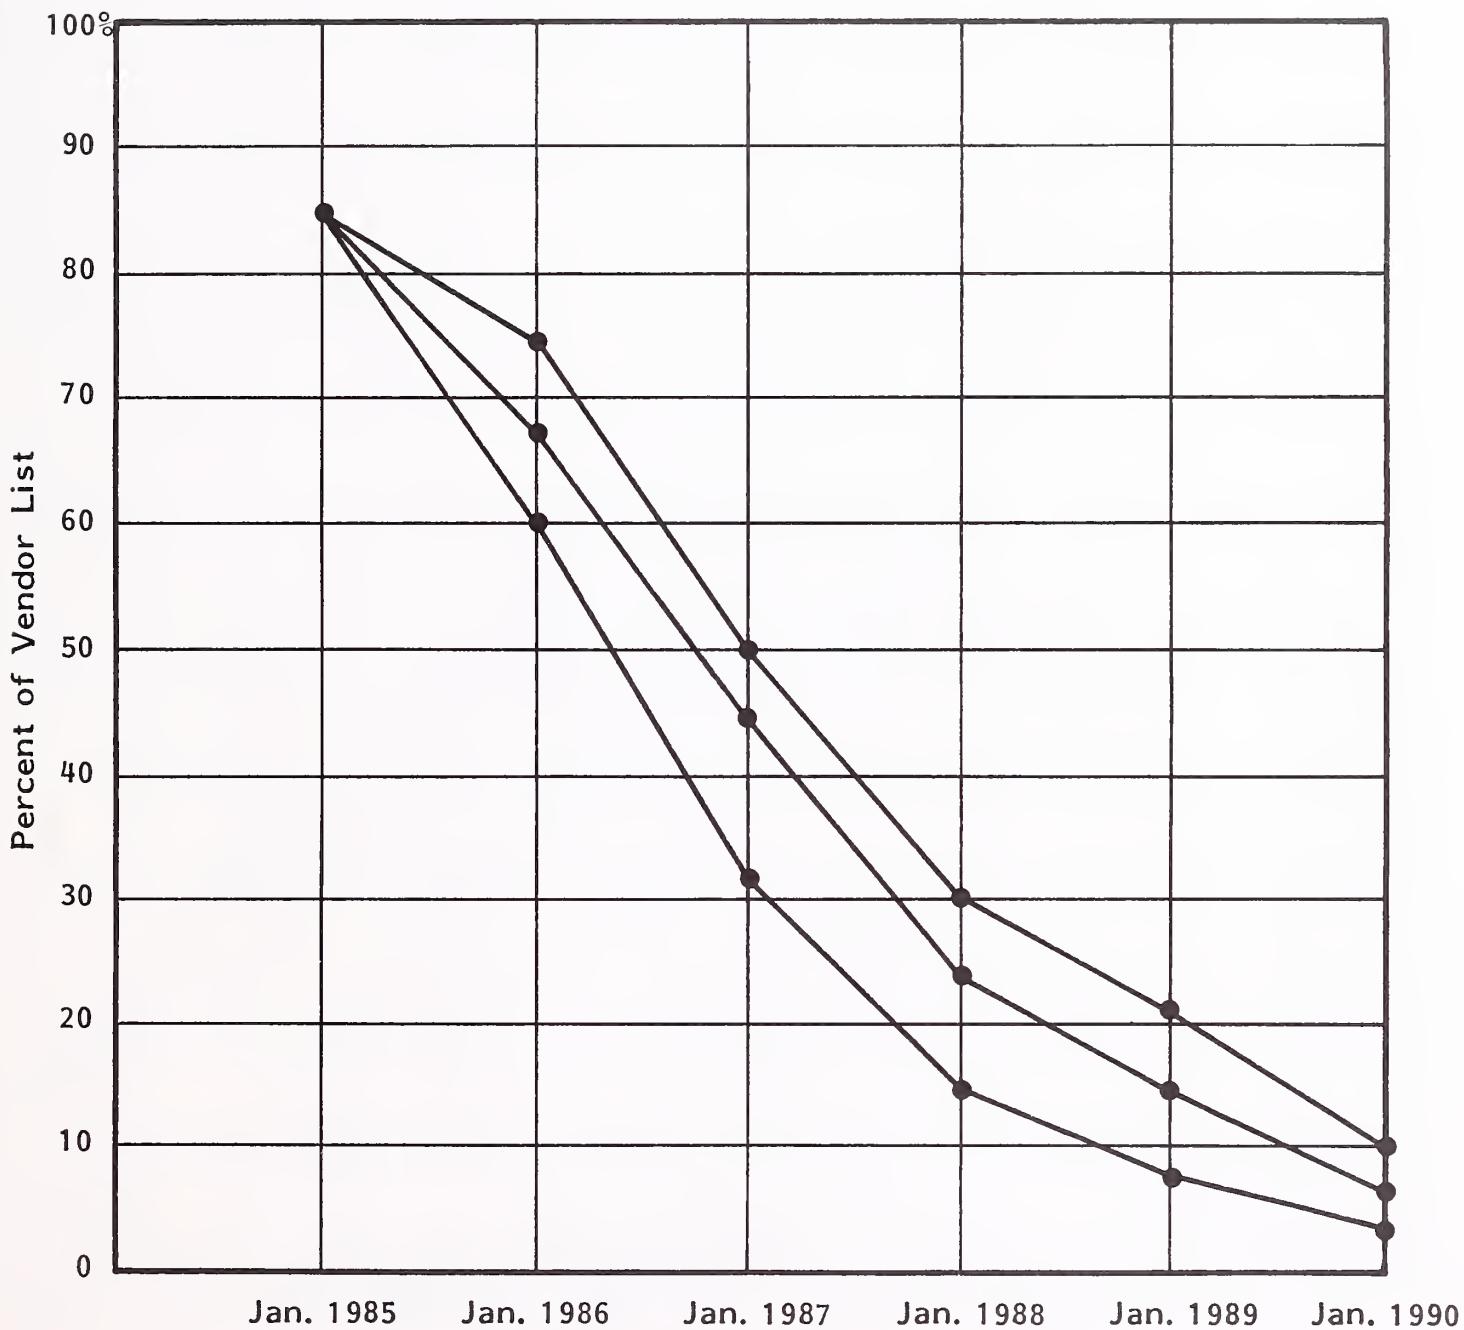
PROJECTED VALUES RANGE	JAN. 1986	JAN. 1987	JAN. 1988	JAN. 1989	JAN. 1990
High	65%	43%	25%	15%	7%
Expected	60	40	21	10	3
Low	48	30	15	8	1

EXHIBIT III-11
RESIDUAL VALUE FORECAST FOR
IBM 3081-GX16 PROCESSOR



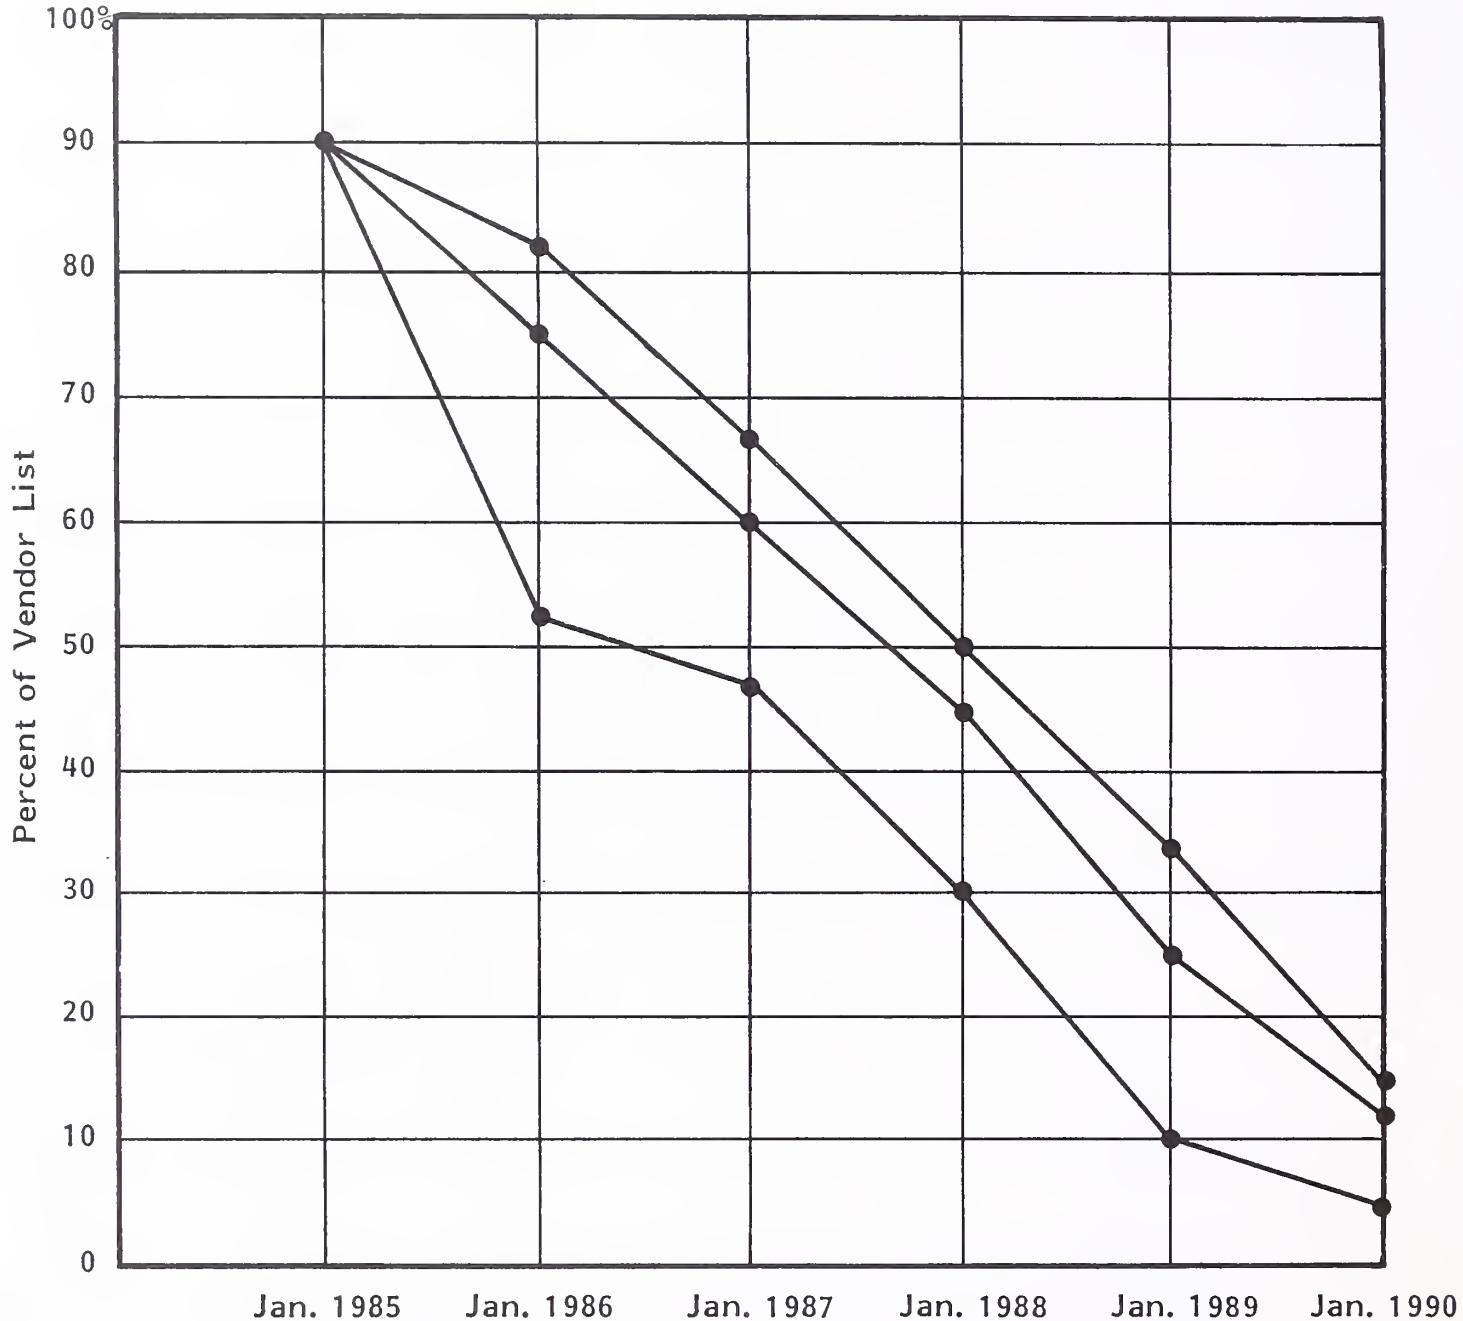
PROJECTED VALUES RANGE	JAN. 1986	JAN. 1987	JAN. 1988	JAN. 1989	JAN. 1990
High	70%	48%	26%	17%	9%
Expected	65	40	20	12	5
Low	52	30	14	7	3

EXHIBIT III-12
RESIDUAL VALUE FORECAST FOR
IBM 3081-KX24 PROCESSOR



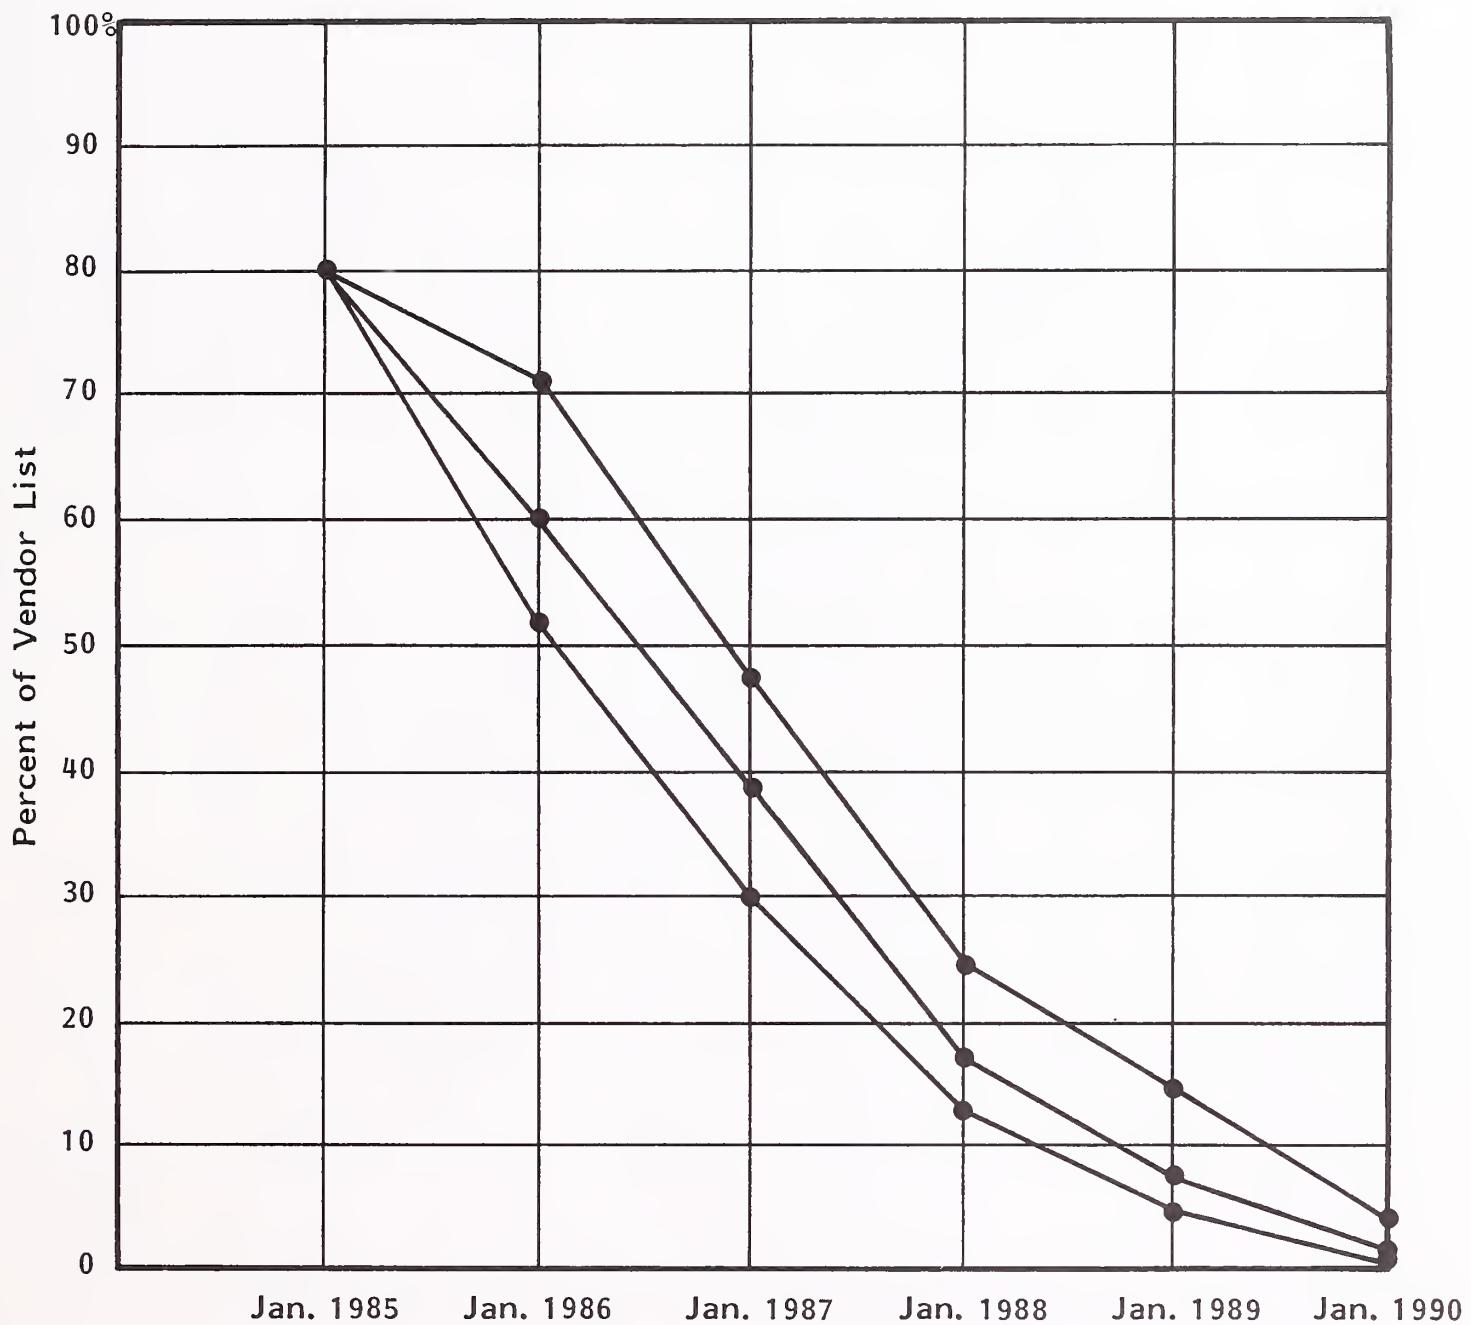
PROJECTED VALUES RANGE	JAN. 1986	JAN. 1987	JAN. 1988	JAN. 1989	JAN. 1990
High	75%	50%	30%	21%	10%
Expected	68	45	24	15	7
Low	60	32	15	8	4

EXHIBIT III-13
RESIDUAL VALUE FORECAST FOR
IBM 3084-QX64 PROCESSOR



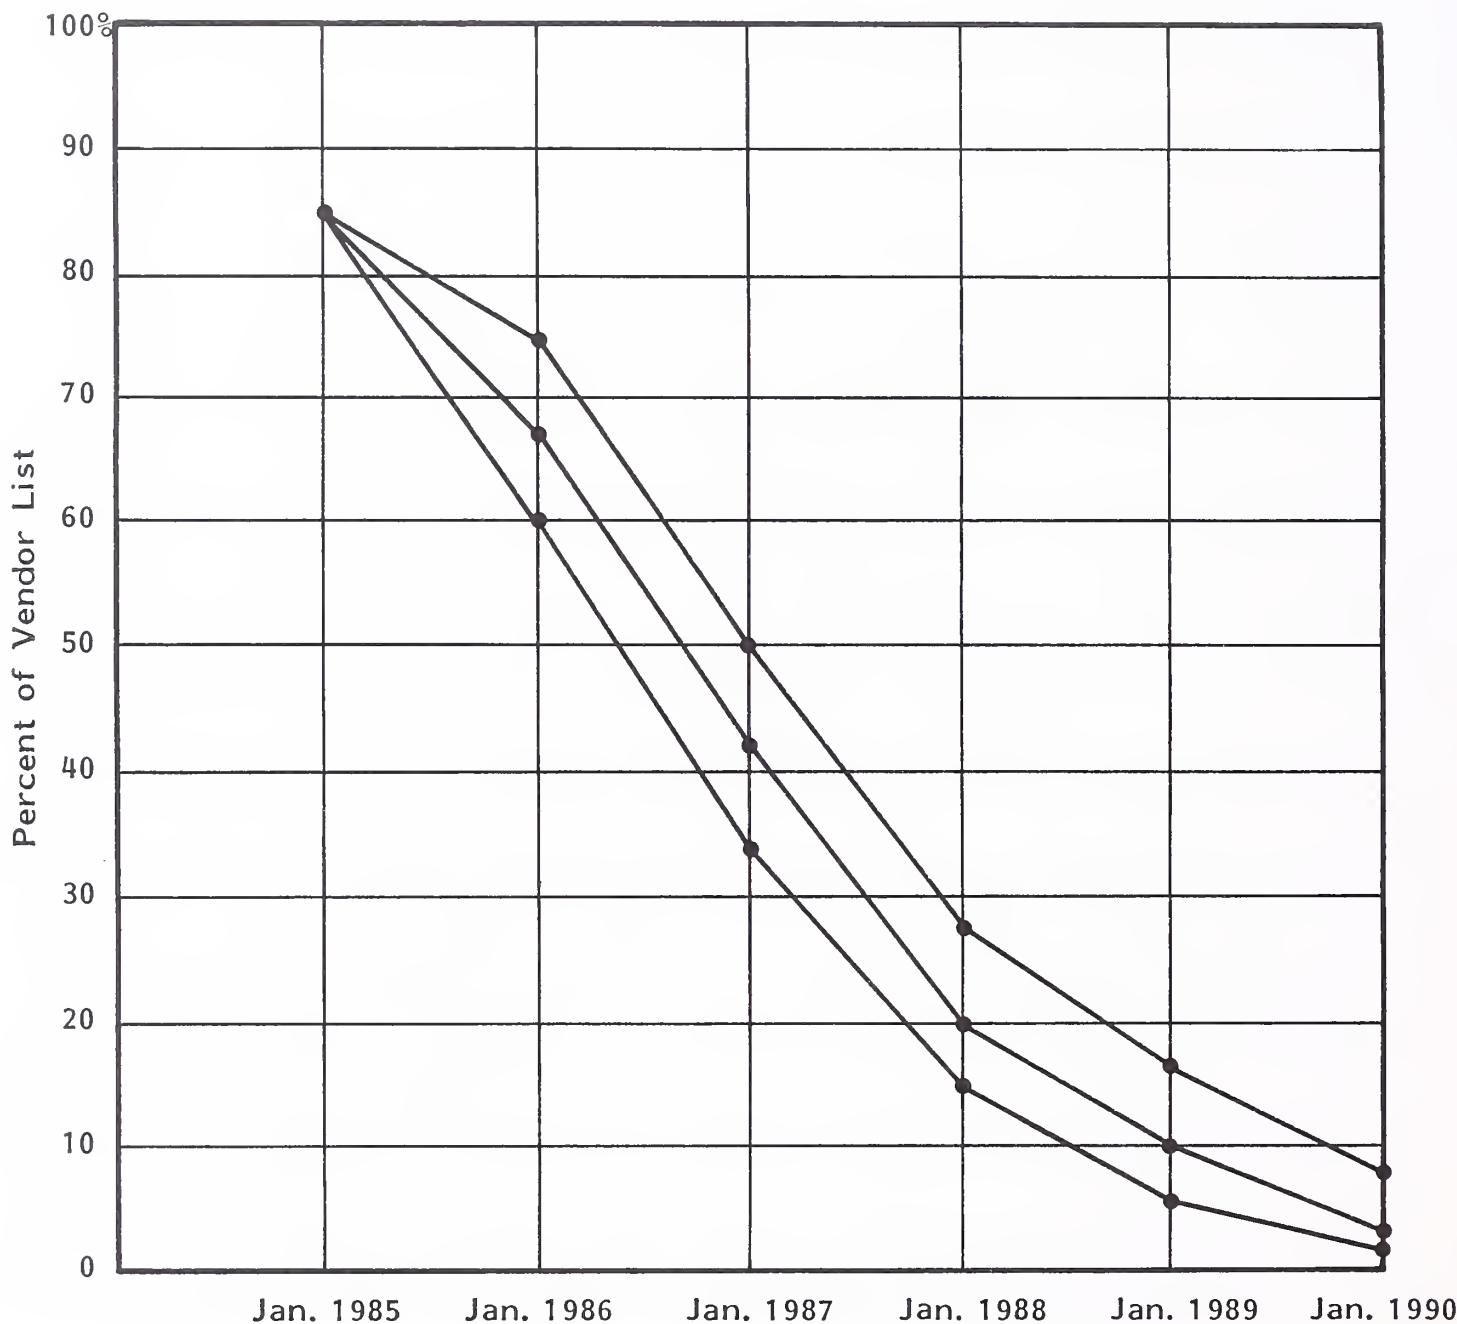
PROJECTED VALUES RANGE	JAN. 1986	JAN. 1987	JAN. 1988	JAN. 1989	JAN. 1990
High	82%	67%	50%	33%	15%
Expected	75	60	45	25	12
Low	52	48	30	10	5

EXHIBIT III-14
RESIDUAL VALUE FORECAST FOR
AMDAHL 5850-24 PROCESSOR



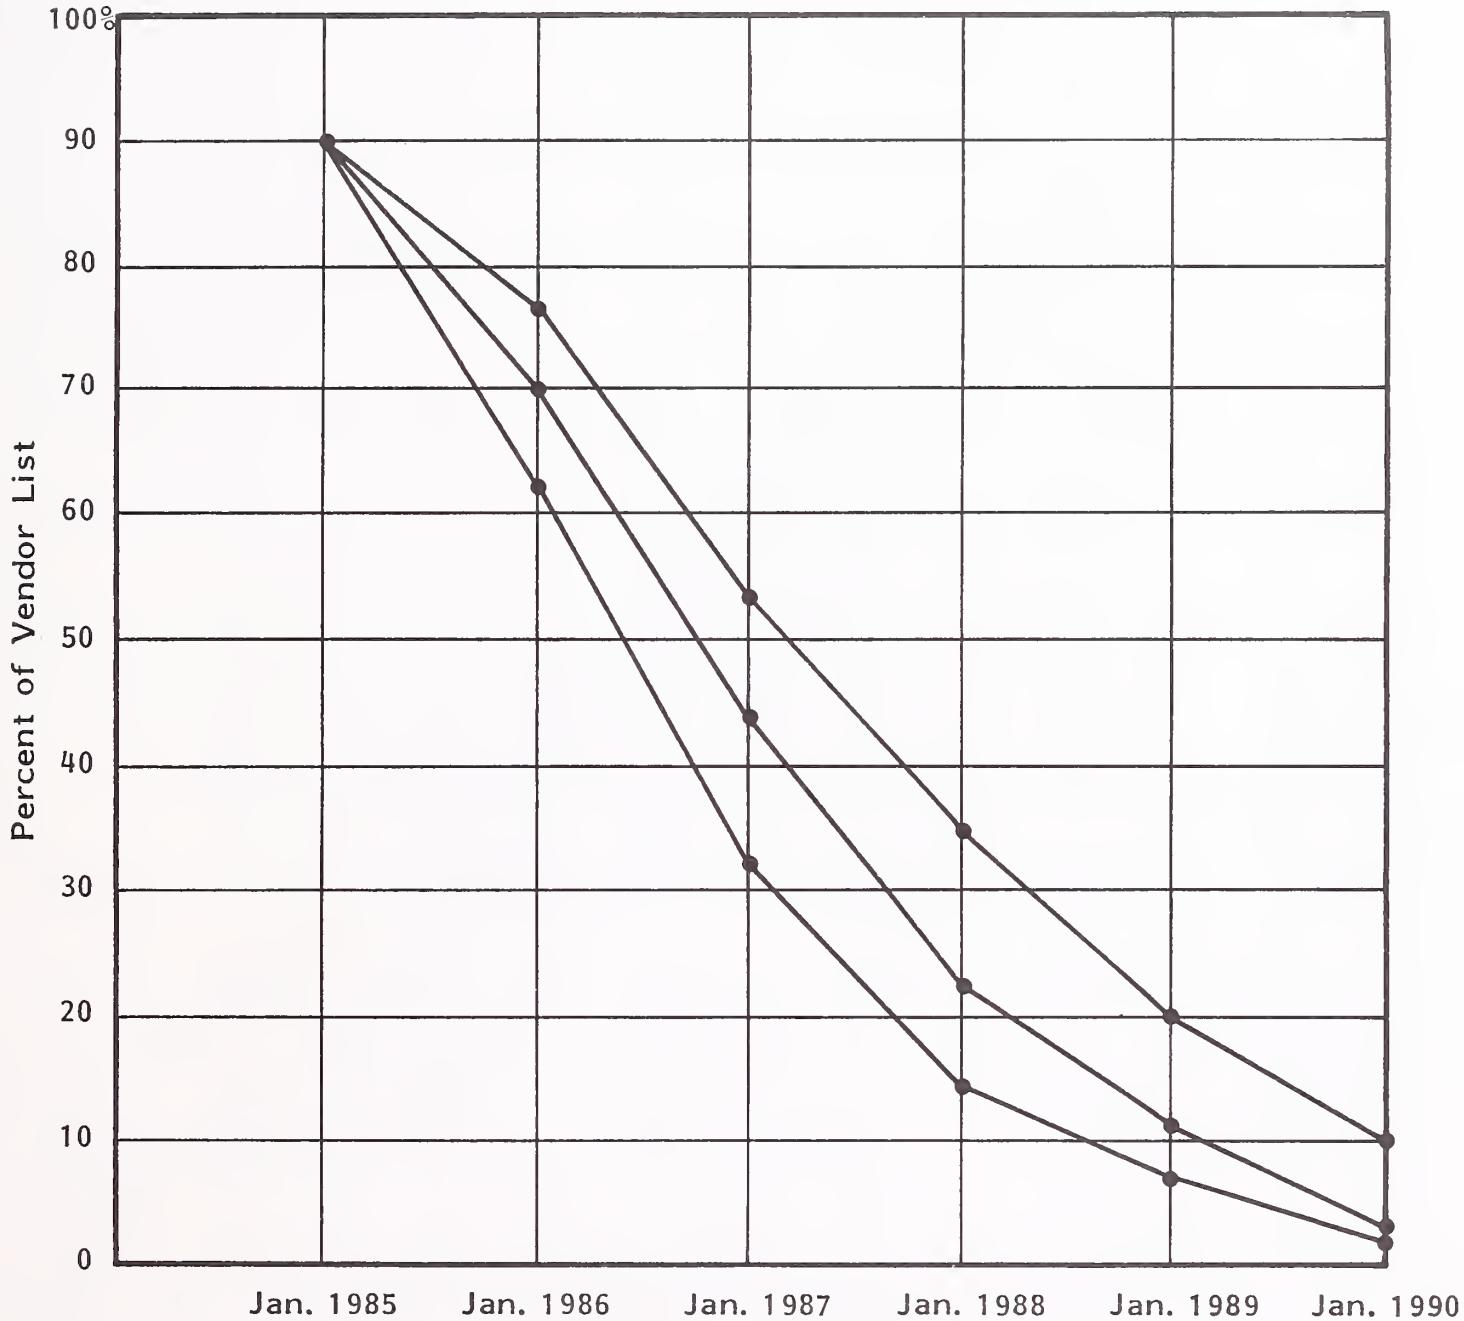
PROJECTED VALUES RANGE	JAN. 1986	JAN. 1987	JAN. 1988	JAN. 1989	JAN. 1990
High	71%	48%	25%	15%	4%
Expected	60	39	18	8	2
Low	52	30	13	5	1

EXHIBIT III-15
RESIDUAL VALUE FORECAST FOR
AMDAHL 5860-24 PROCESSOR



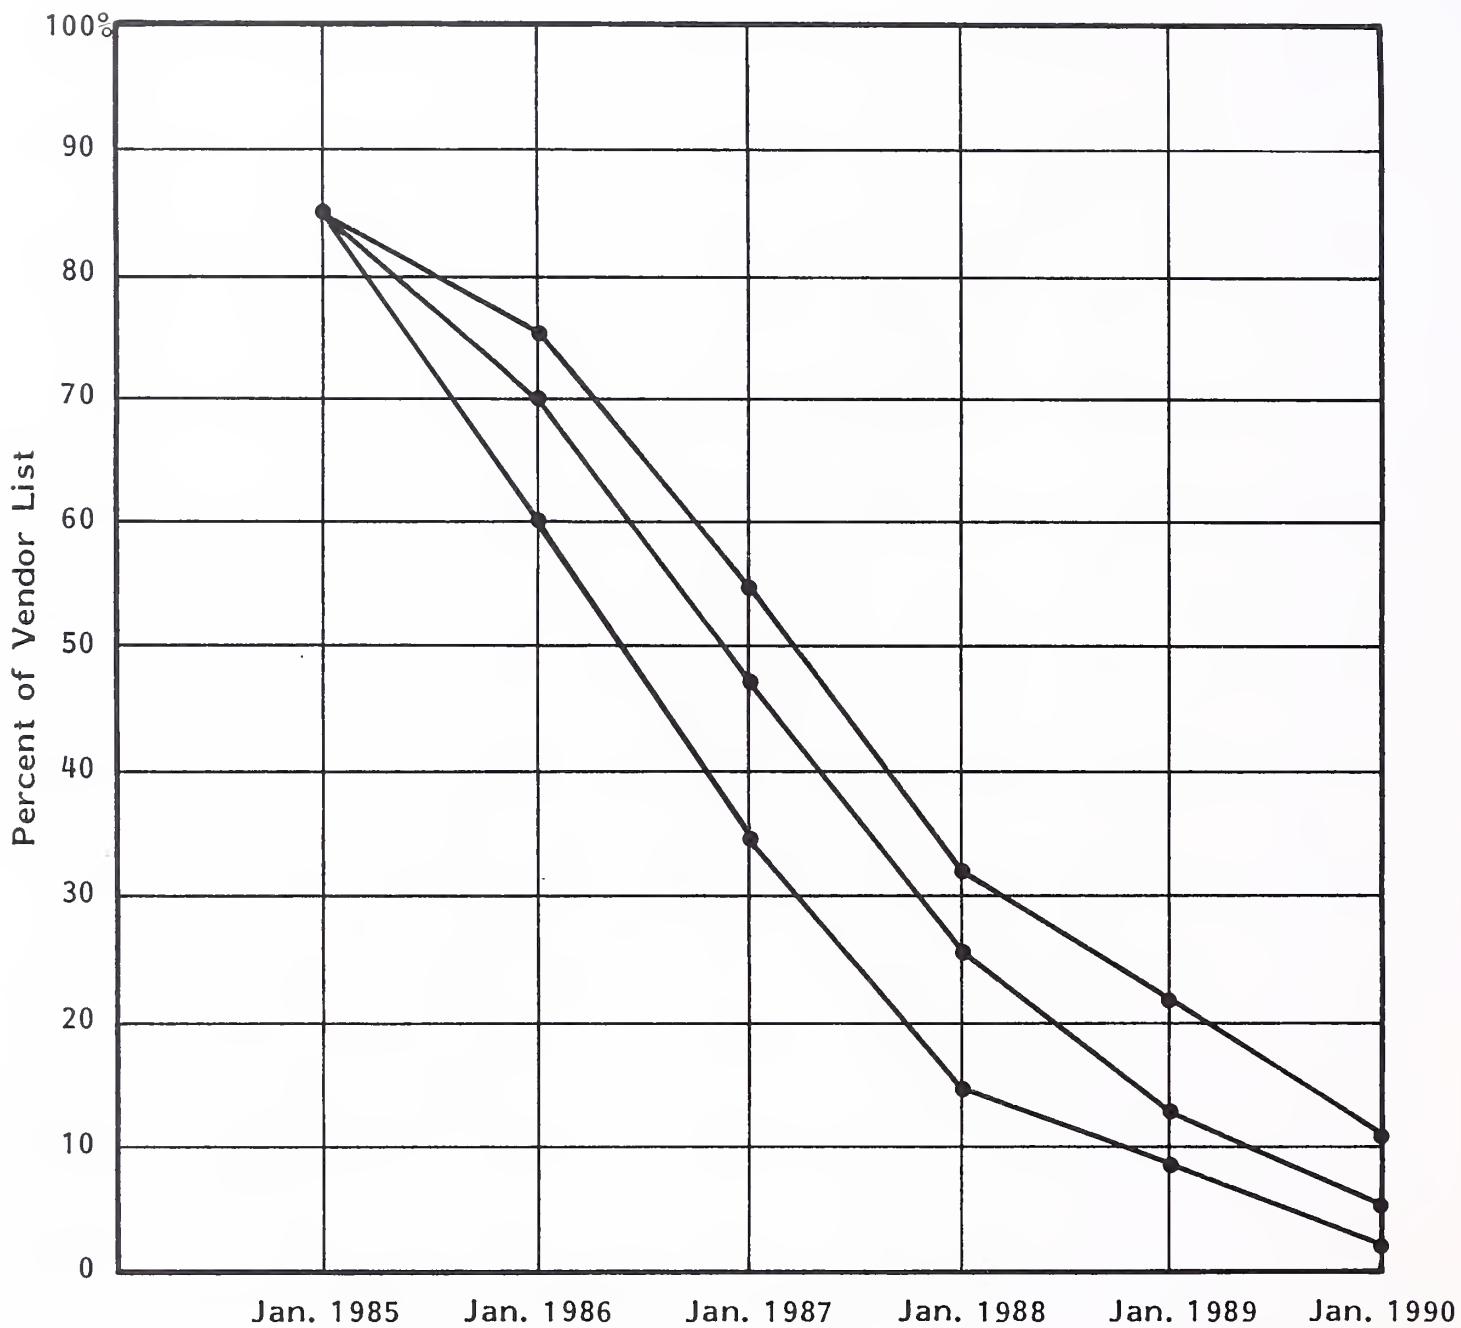
PROJECTED VALUES RANGE	JAN. 1986	JAN. 1987	JAN. 1988	JAN. 1989	JAN. 1990
High	75%	50%	28%	17%	8%
Expected	68	42	20	10	3
Low	60	34	15	6	1

EXHIBIT III-16
RESIDUAL VALUE FORECAST FOR
AMDAHL 5868-32 PROCESSOR



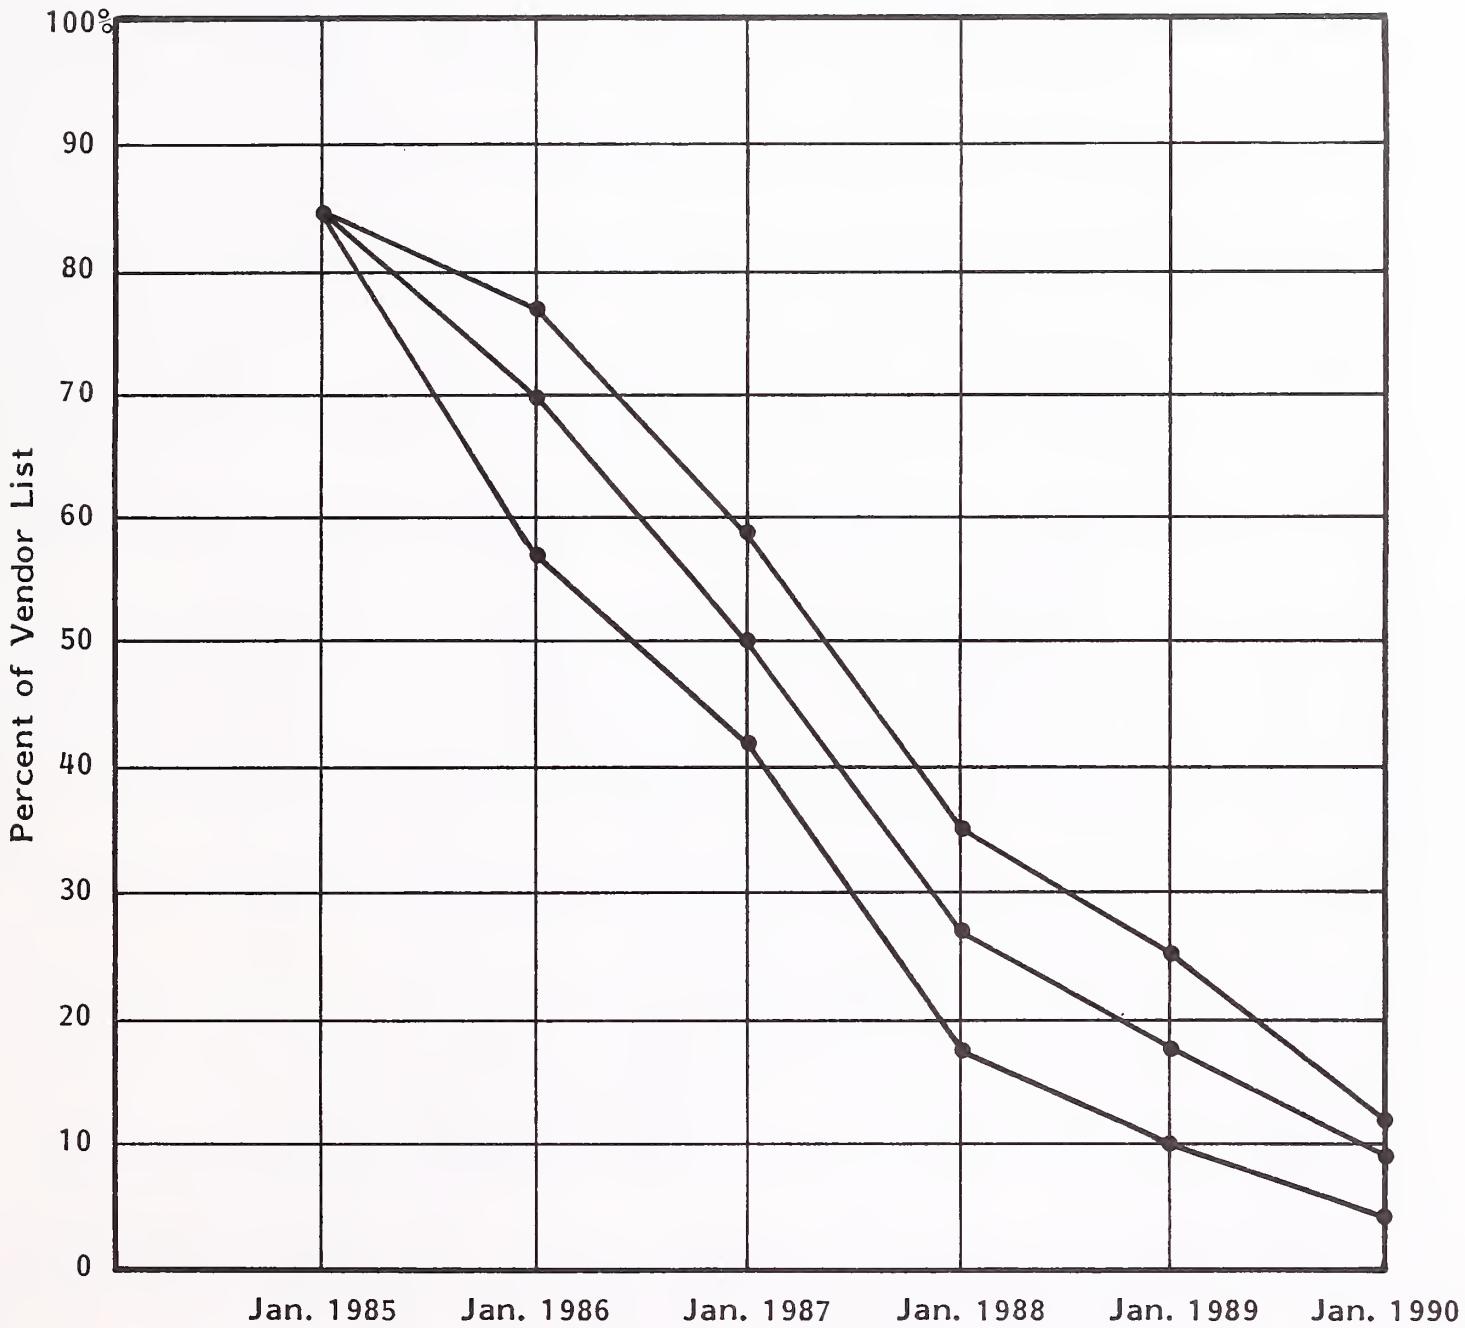
PROJECTED VALUES RANGE	JAN. 1986	JAN. 1987	JAN. 1988	JAN. 1989	JAN. 1990
High	77%	53%	35%	20%	10%
Expected	70	44	22	11	3
Low	62	32	15	8	2

EXHIBIT III-17
RESIDUAL VALUE FORECAST FOR
AMDAHL 5870-32 PROCESSOR



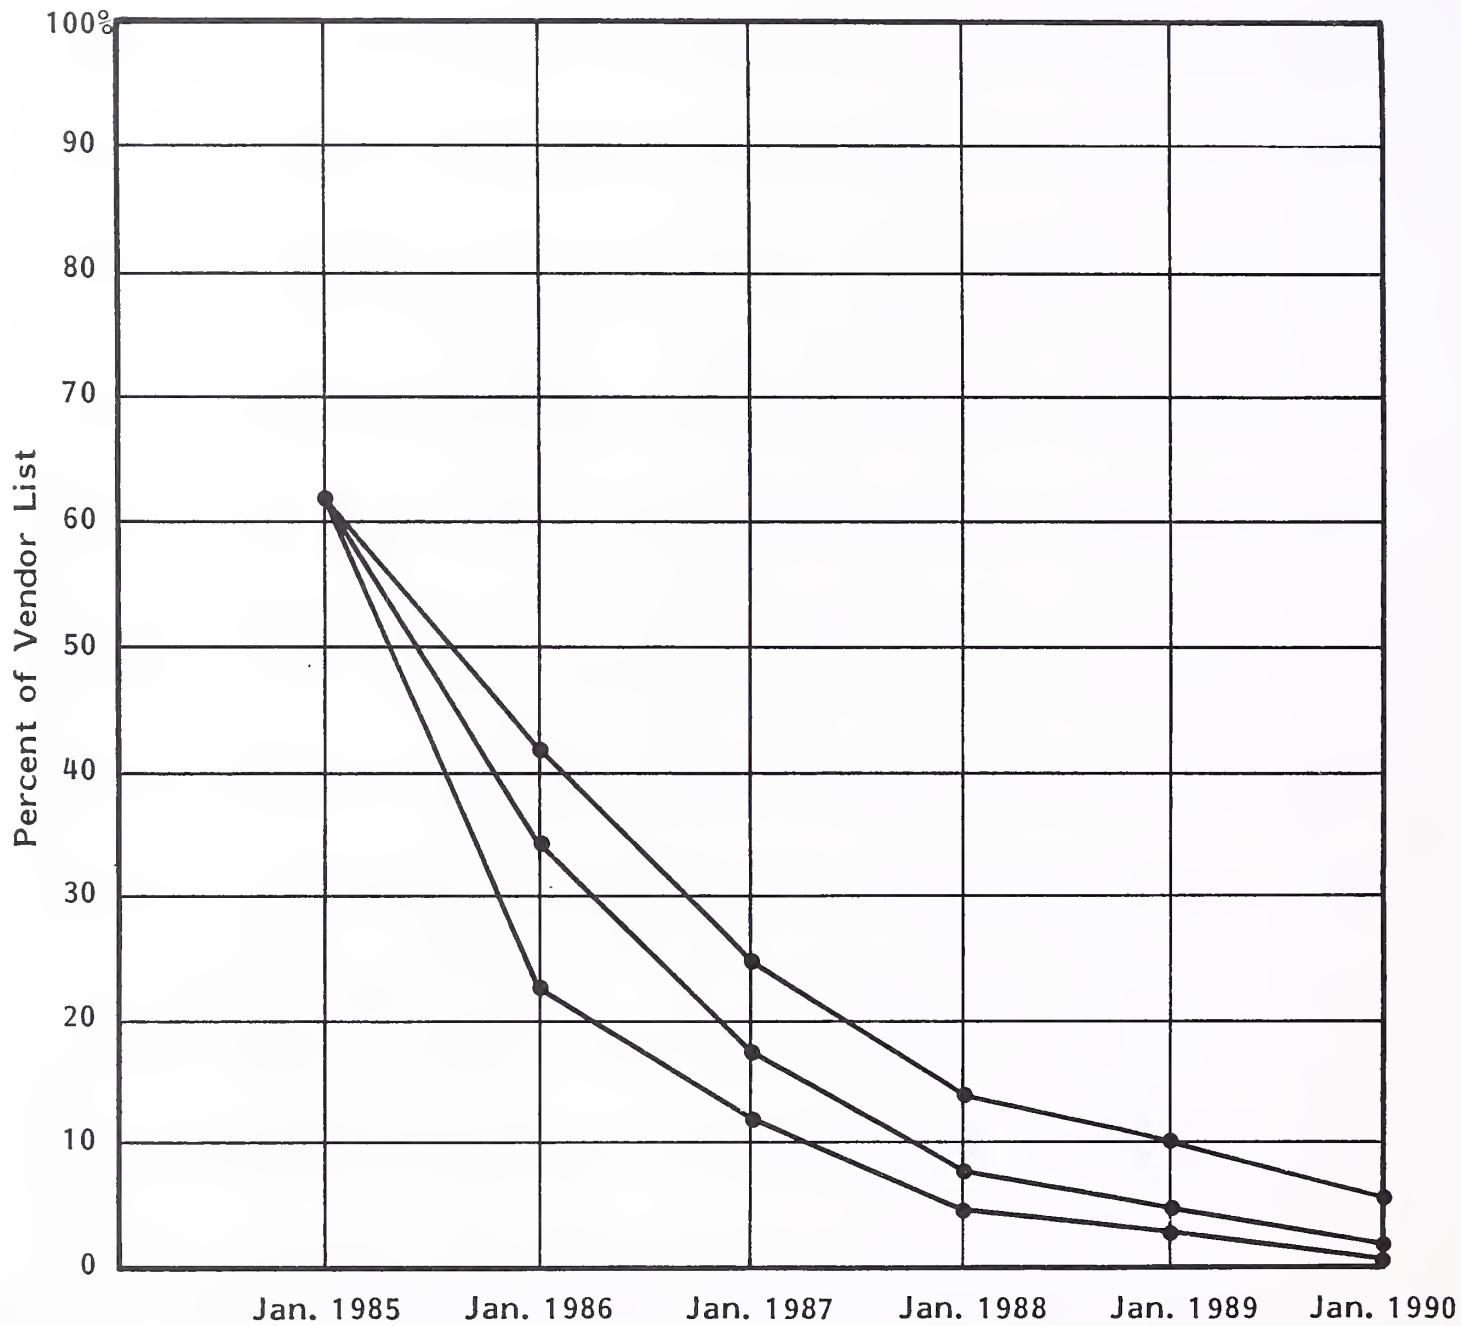
PROJECTED VALUES RANGE	JAN. 1986	JAN. 1987	JAN. 1988	JAN. 1989	JAN. 1990
High	76%	55%	32%	22%	11%
Expected	70	48	26	13	6
Low	60	35	15	9	2

EXHIBIT III-18
RESIDUAL VALUE FORECAST FOR
AMDAHL 5880-64 PROCESSOR



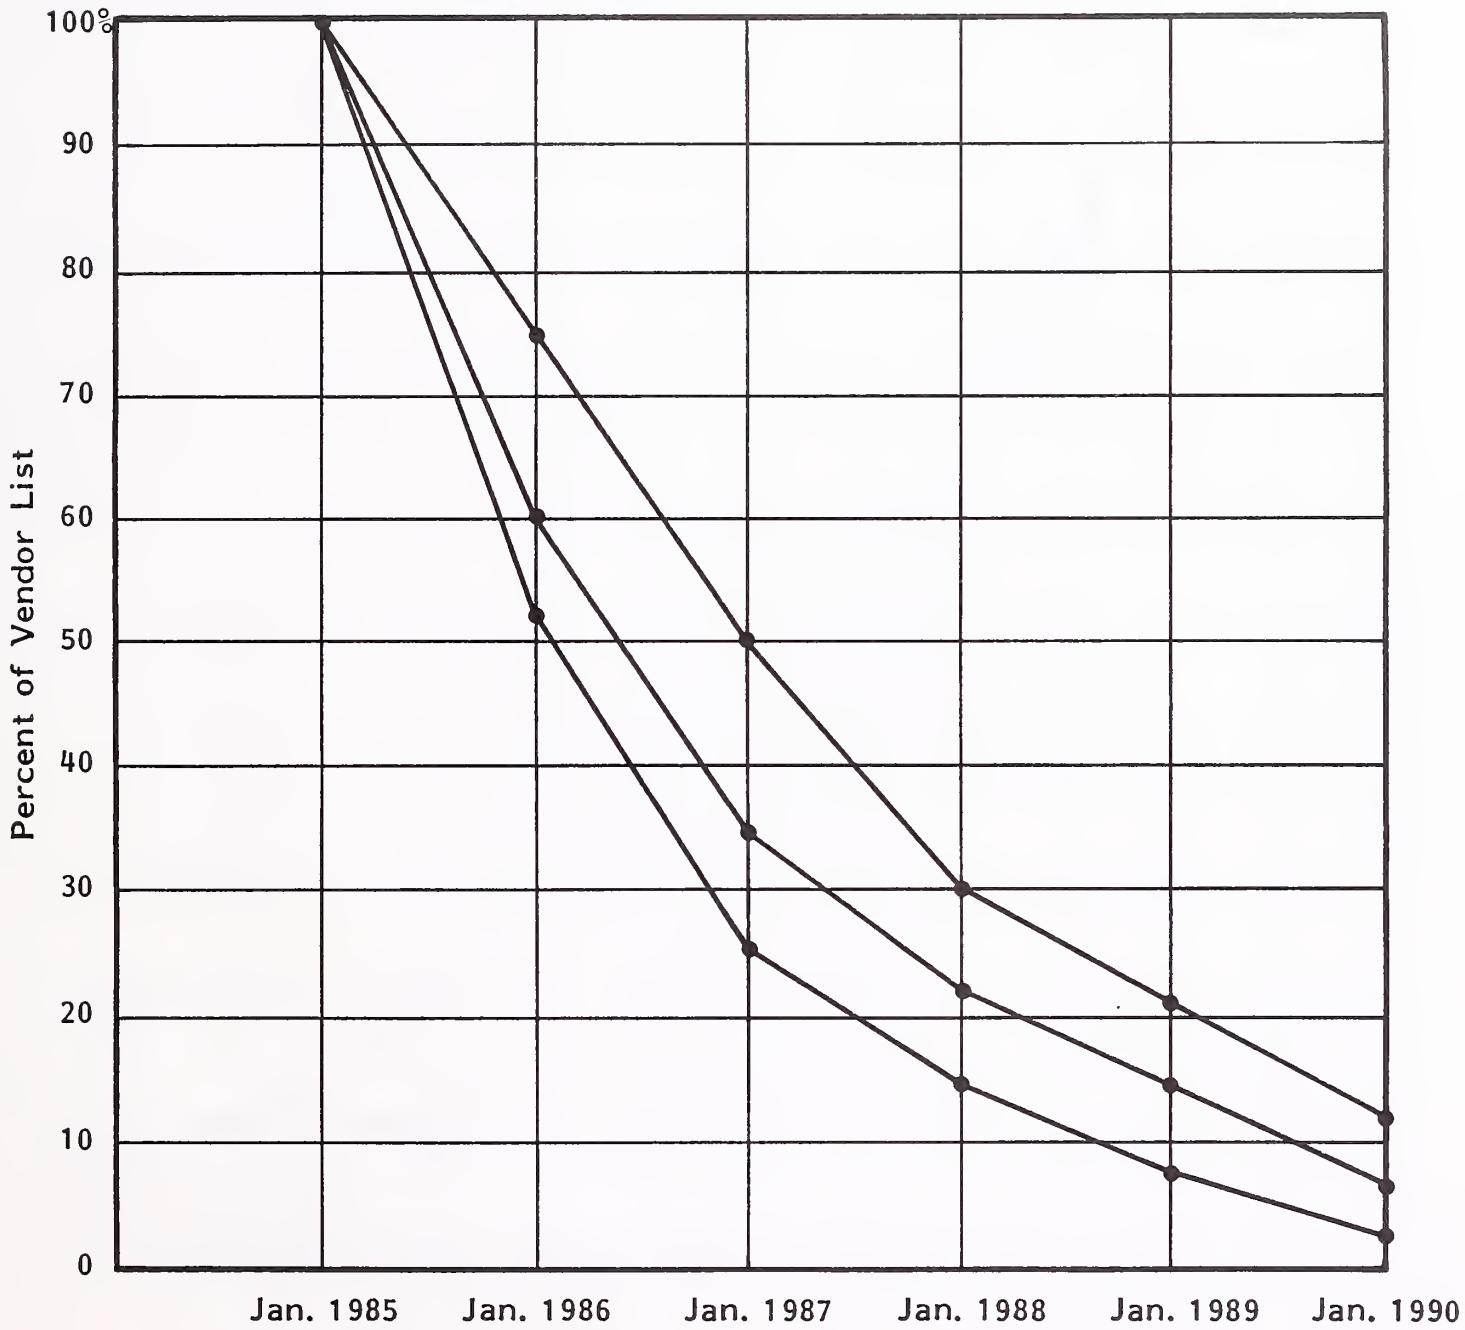
PROJECTED VALUES RANGE	JAN. 1986	JAN. 1987	JAN. 1988	JAN. 1989	JAN. 1990
High	78%	59%	35%	25%	12%
Expected	70	50	28	18	9
Low	58	42	18	10	4

EXHIBIT III-19
RESIDUAL VALUE FORECAST FOR
NAS AS/6630



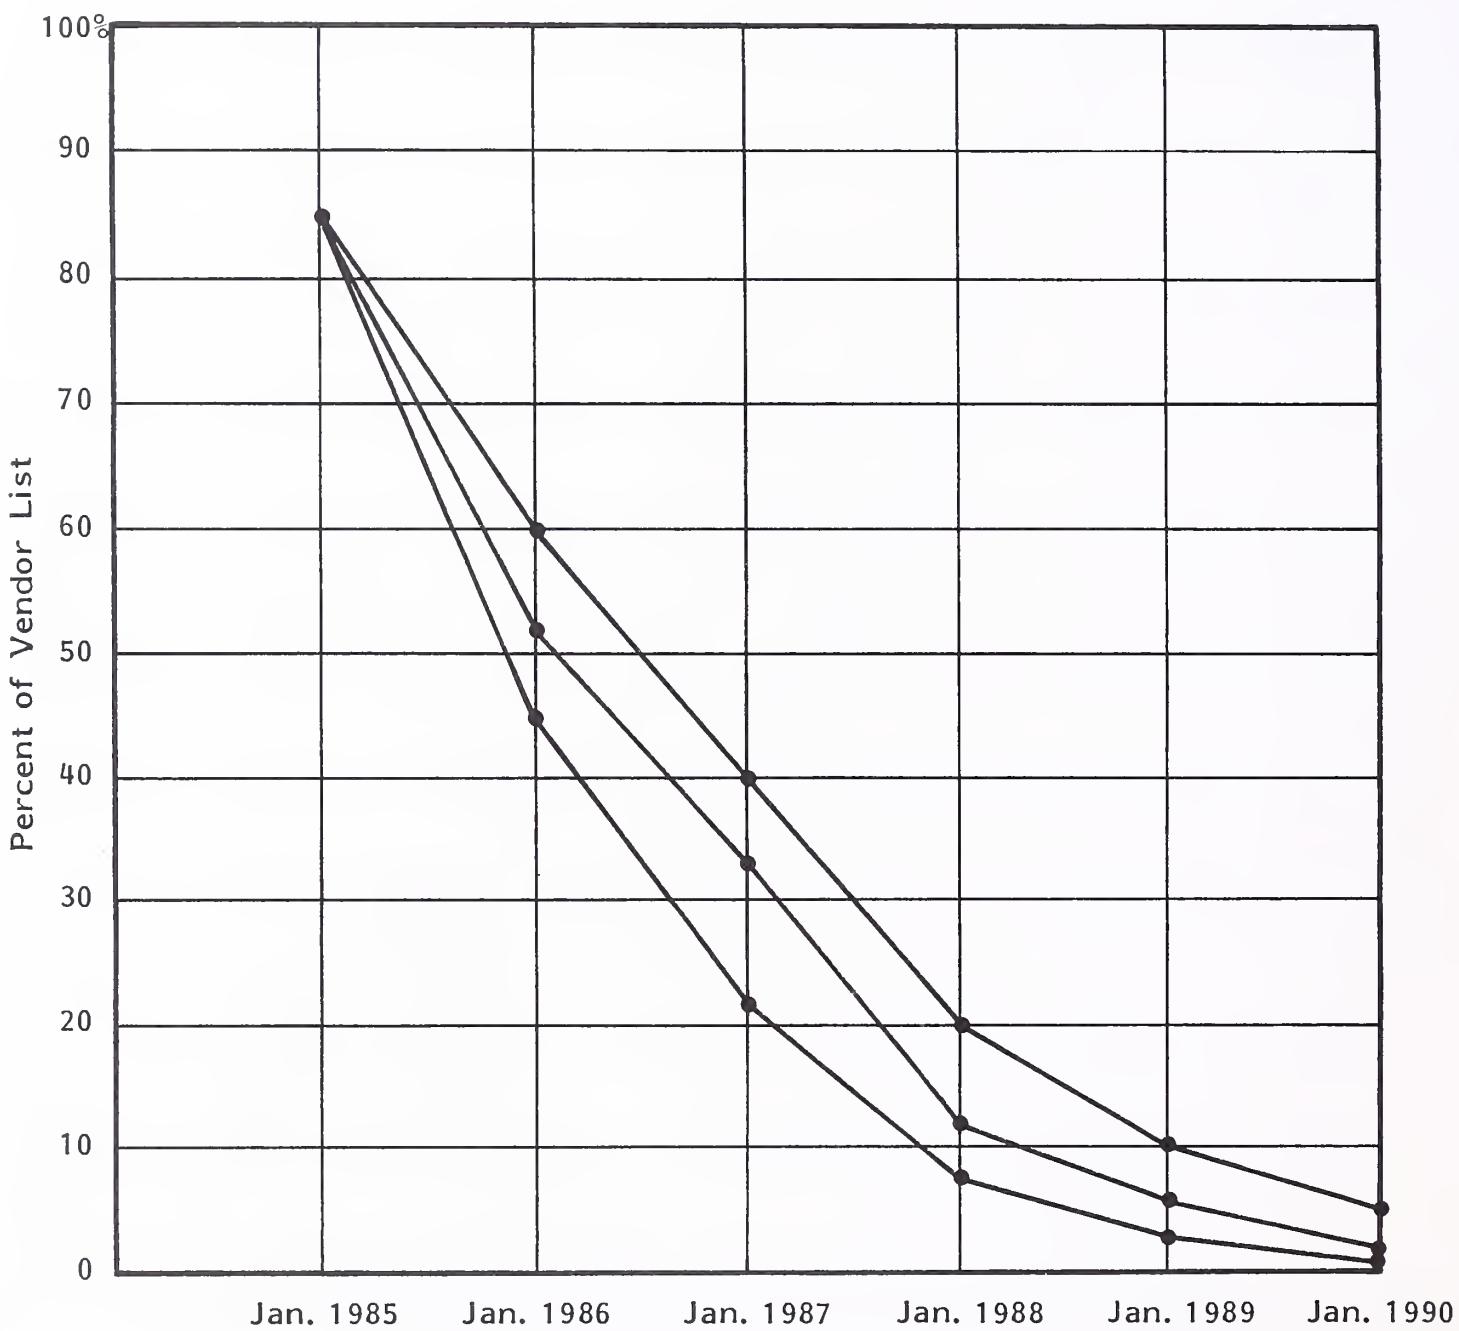
PROJECTED VALUES RANGE	JAN. 1986	JAN. 1987	JAN. 1988	JAN. 1989	JAN. 1990
High	42%	25%	14%	10%	6%
Expected	35	18	8	5	2
Low	23	12	5	3	1

EXHIBIT III-20
RESIDUAL VALUE FORECAST FOR
NAS AS/6660



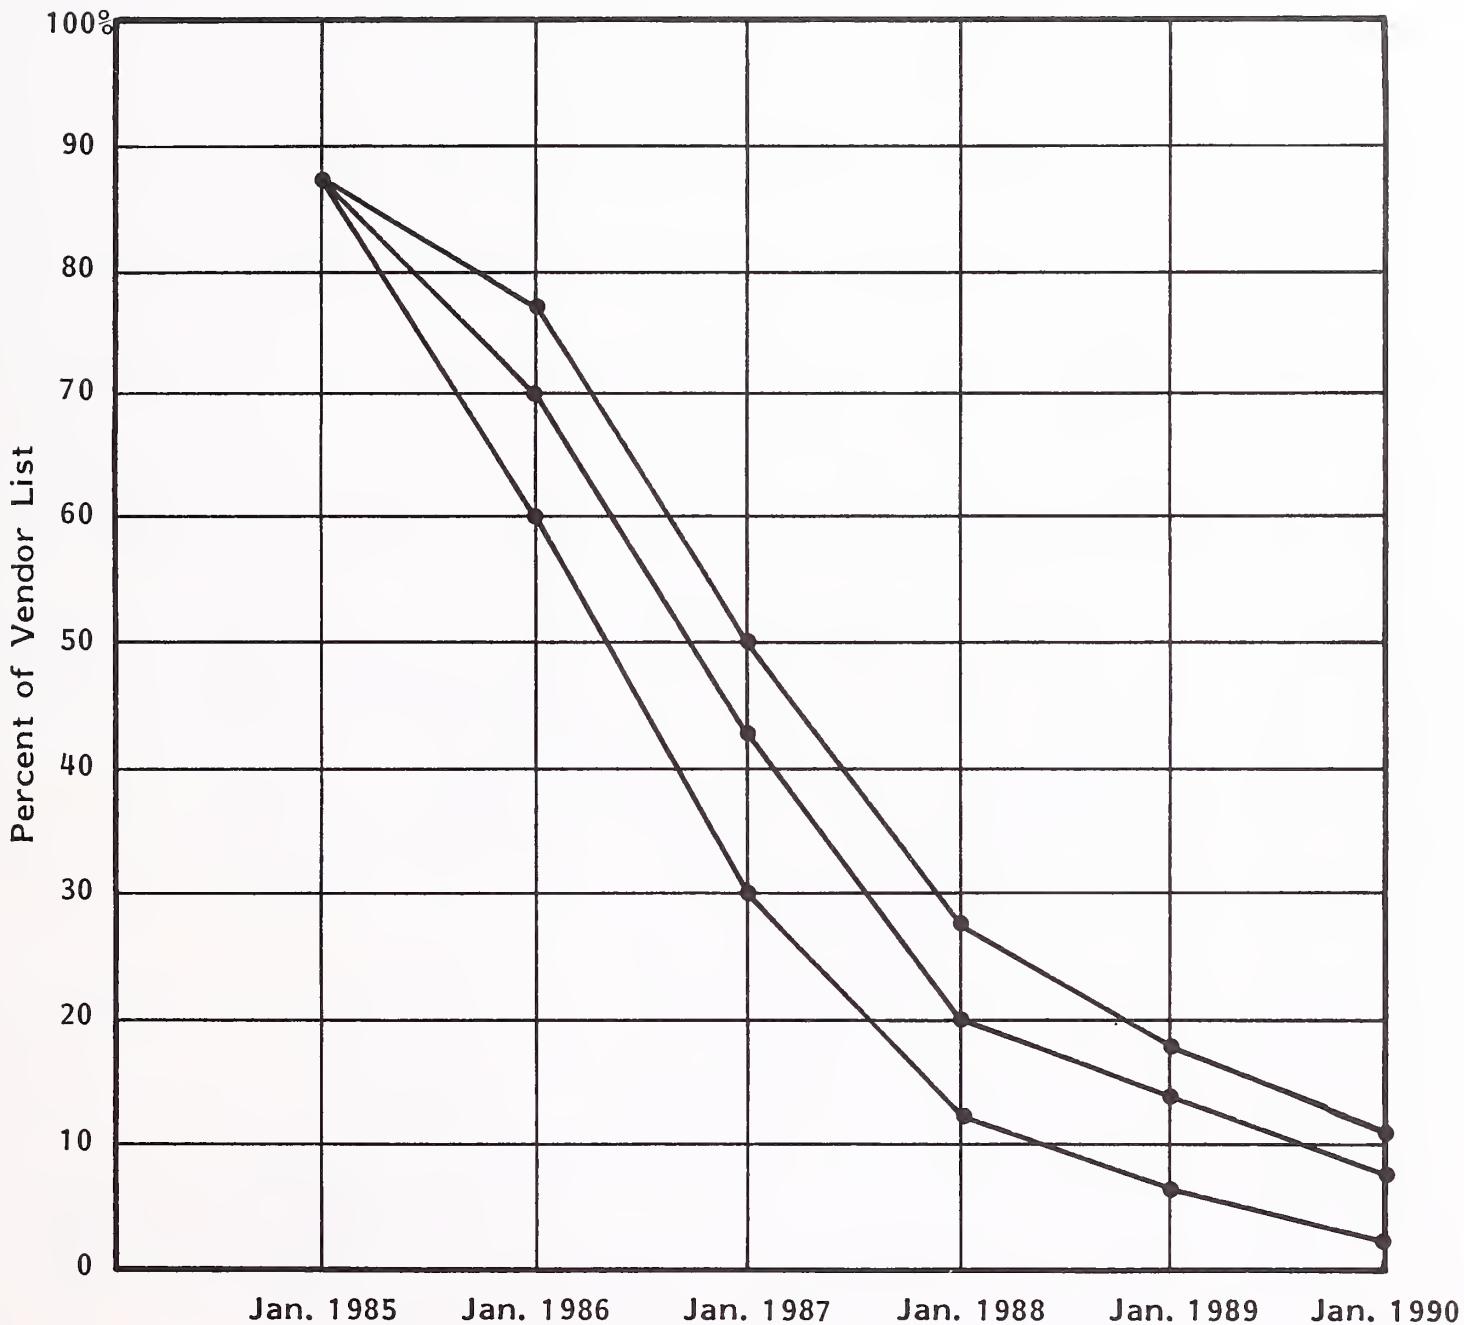
PROJECTED VALUES RANGE	JAN. 1986	JAN. 1987	JAN. 1988	JAN. 1989	JAN. 1990
High	75%	50%	30%	21%	12%
Expected	60	35	22	15	7
Low	52	26	15	8	3

EXHIBIT III-21
RESIDUAL VALUE FORECAST FOR
NAS AS/8023



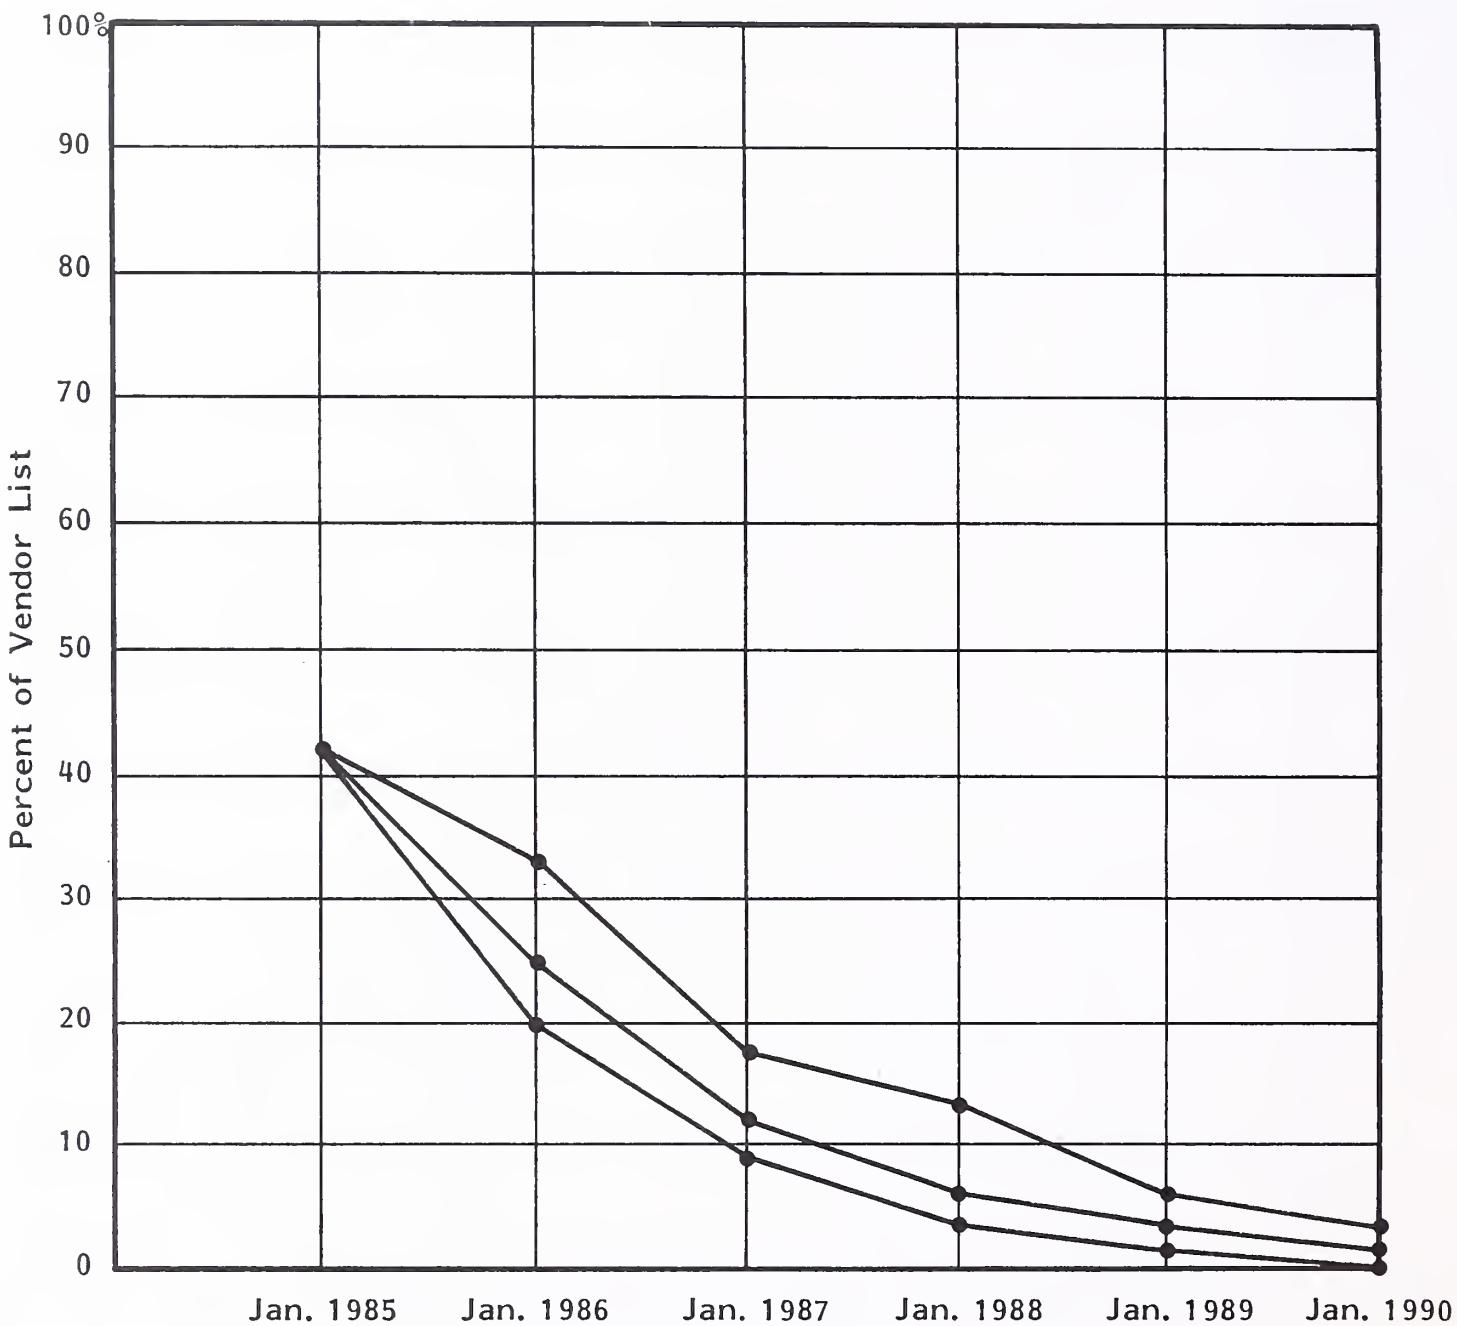
PROJECTED VALUES RANGE	JAN. 1986	JAN. 1987	JAN. 1988	JAN. 1989	JAN. 1990
High	60%	40%	20%	10%	5%
Expected	52	33	12	6	2
Low	45	22	8	3	1

EXHIBIT III-22
RESIDUAL VALUE FORECAST FOR
NAS AS/8083



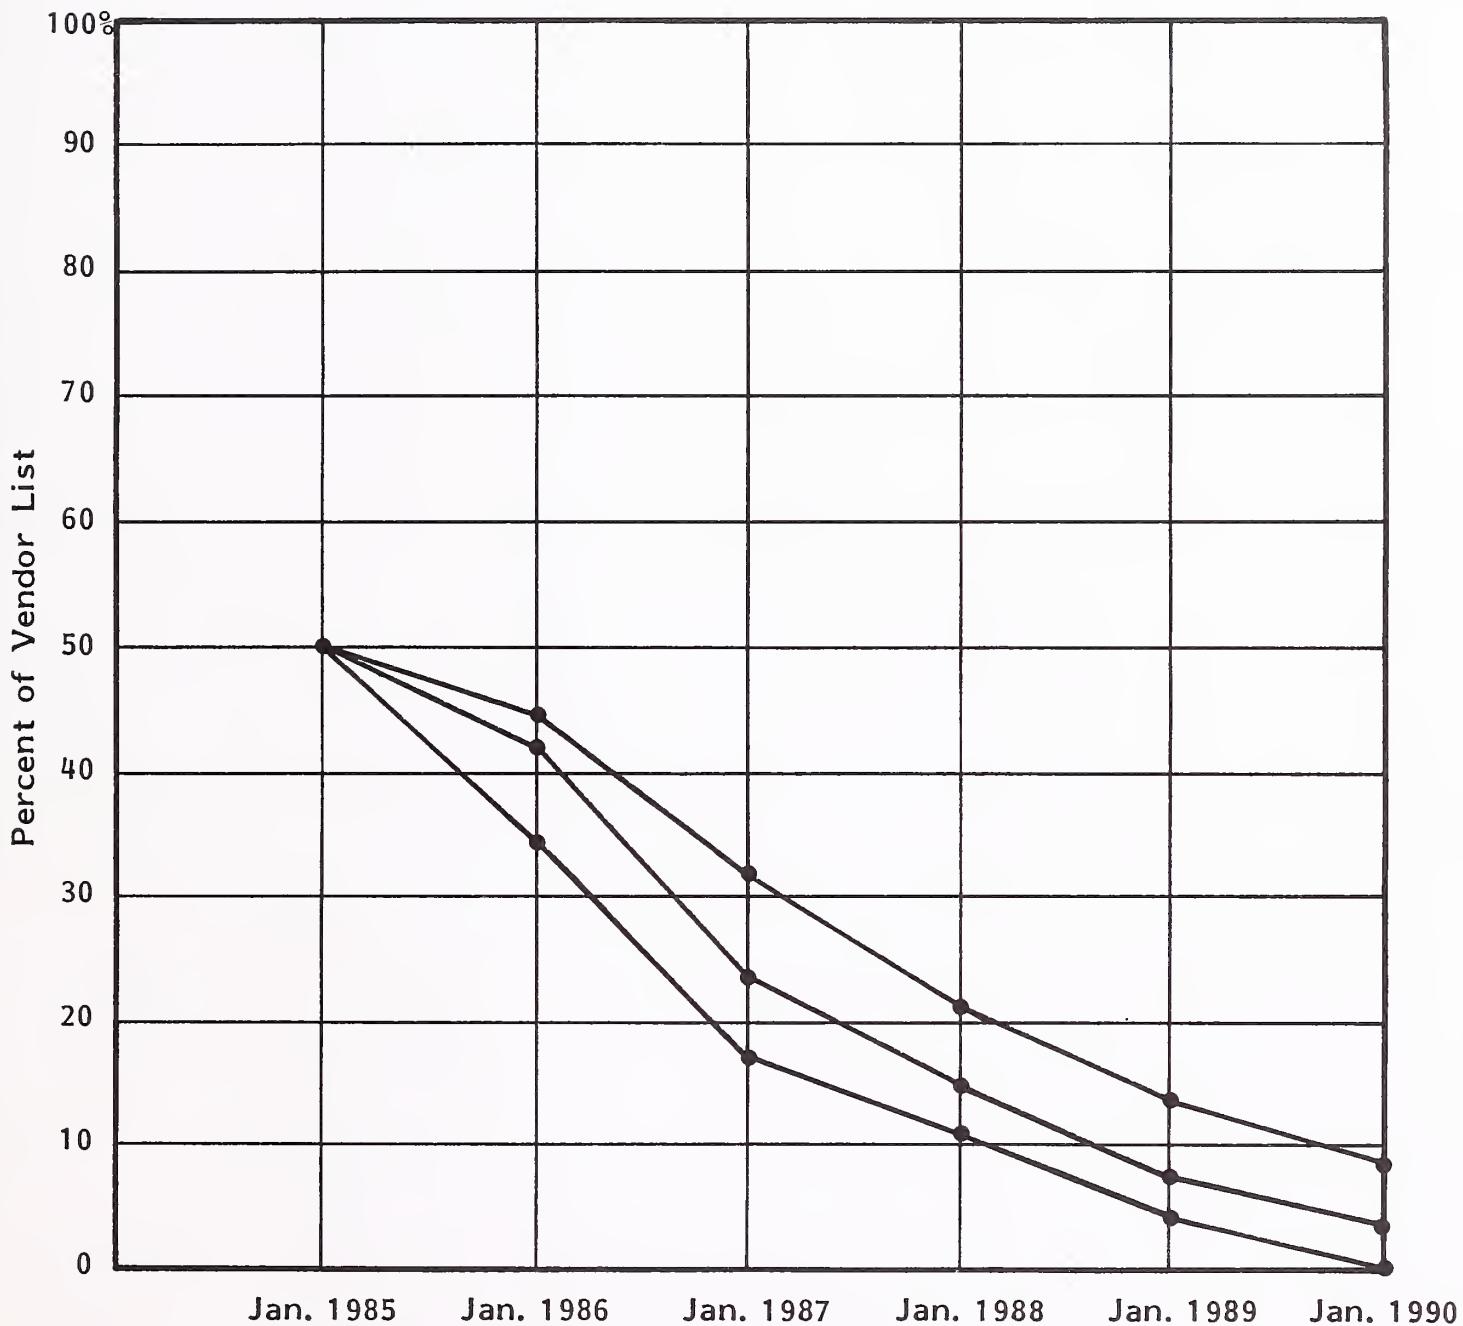
PROJECTED VALUES RANGE	JAN. 1986	JAN. 1987	JAN. 1988	JAN. 1989	JAN. 1990
High	78%	50%	28%	18%	11%
Expected	70	43	20	14	8
Low	60	30	12	7	2

EXHIBIT III-23
RESIDUAL VALUE FORECAST FOR
NAS AS/9050



PROJECTED VALUES RANGE	JAN. 1986	JAN. 1987	JAN. 1988	JAN. 1989	JAN. 1990
High	33%	18%	13%	7%	4%
Expected	25	12	7	4	2
Low	20	9	4	2	0

EXHIBIT III-24
RESIDUAL VALUE FORECAST FOR
NAS AS/9070



PROJECTED VALUES RANGE	JAN. 1986	JAN. 1987	JAN. 1988	JAN. 1989	JAN. 1990
High	45%	32%	21%	14%	9%
Expected	42	24	15	8	4
Low	35	18	11	5	0

INPUT[®] 1943 Landings Drive, Mountain View, CA 94043, (415) 960-3990